

Information technology - SCSI Object-Based Storage Device Commands -2 (OSD-2)

This is an internal working document of T10, a Technical Committee of Accredited Standards Committee INCITS (InterNational Committee for Information Technology Standards). As such this is not a completed standard and has not been approved. The contents may be modified by the T10 Technical Committee. The contents are actively being modified by T10. This document is made available for review and comment only.

Permission is granted to members of INCITS, its technical committees, and their associated task groups to reproduce this document for the purposes of INCITS standardization activities without further permission, provided this notice is included. All other rights are reserved. Any duplication of this document for commercial or for-profit use is strictly prohibited.

T10 Technical Editor:

Ralph O. Weber
ENDL Texas
18484 Preston Road
Suite 102 PMB 178
Dallas, TX 75252
USA

Telephone: 214-912-1373
Facsimile: 972-596-2775
Email: ROWeber at IEEE dot org

Funded by Seagate Technology

Reference number
ISO/IEC 14776-392 : 200x
ANSI INCITS.***:200x

Points of Contact:

T10 Chair

John B. Lohmeyer
LSI Corp.
4420 Arrows West Drive
Colorado Springs, CO 80907-3444
Tel: (719) 533-7560
Fax: (719) 533-7183
Email: lohmeier at t10 dot org

T10 Vice-Chair

Mark Evans
Western Digital
5863 Rue Ferrari
San Jose, CA 95138
Tel: (408) 363-5257
Fax: (408) 363-5139
Email: Mark dot Evans at wdc dot com

INCITS Secretariat

INCITS Secretariat
1250 Eye Street, NW Suite 200
Washington, DC 20005

Telephone: 202-737-8888
Facsimile: 202-638-4922
Email: incits@itic.org

T10 Web Site www.t10.org

T10 Reflector To subscribe send e-mail to majordomo@T10.org with 'subscribe' in message body
To unsubscribe send e-mail to majordomo@T10.org with 'unsubscribe' in message body

Document Distribution

INCITS Online Store
managed by Techstreet
1327 Jones Drive
Ann Arbor, MI 48105

<http://www.techstreet.com/incits.html>
Telephone: 1-734-302-7801 or
1-800-699-9277
Facsimile: 1-734-302-7811

or

Global Engineering
15 Inverness Way East
Englewood, CO 80112-5704

<http://global.ihs.com/>
Telephone: 1-303-792-2181 or
1-800-854-7179
Facsimile: 1-303-792-2192

Revision Information

1 Approved Documents Included

The following T10 approved proposals have been incorporated OSD-2 up to and including this revision:

04-290r1	Condition and event definitions
04-296r1	Numbering conventions
04-313r1	Mandatory REPORT LUNS Support
05-257r0	Agreed editorial corrections
05-311r0	REMOVE PARTITION Security Controlled by Root Attributes
05-314r1	Possible inconsistencies in integrity check value algorithm field definition/usage
05-316r1	Multi-Object LIST and LIST COLLECTION command enhancements
05-328r1	Four New Multi-Object Commands
06-257r0	Remove PREVENT ALLOW MEDIUM REMOVAL
06-259r2	Making linked commands obsolete
07-257r1	Prohibited needed as a keyword in SPC-4
07-270r2	Several OSD-2 Corrections and Clarifications
07-273r3	Attributes Enhancements
07-274r1	CLEAR command, PUNCH command, & range-based FLUSH
07-275r2	Task Management Function Catchup
07-301r5	OSD-2 Security Enhancements
07-357r0	Correct an OSD-2 defined attributes bug
07-378r2	OSD-2 Exceptions Management enhancements
08-041r1	Use period as decimal separator in T10 standards
08-179r3	Fixes for five OSD-2 bugs
08-181r0	Set Attributes error handling in OSD-2
08-182r3	Snapshots and related enhancements
08-185r5	CDB Continuations Definition and Usage

To the best of the technical editor's knowledge, all T10 approved proposals have been included in this revision.

2 Revision History

2.1 Revision 0 (4 October 2004)

Revision 0 of OSD-2 is equal to revision 10 of OSD with the agreed editorial changes shown in:

04-315r0	ANSI Editor Queries on BSR INCITS 400 (OSD)
04-325r0	Editorial corrections for BSR INCITS 400

Revision 0 also adds an informative annex listing the attributes defined by this standard in numerical order.

2.2 Revision 1 (22 January 2007)

The following approved documents have been incorporated in this revision:

04-290r1	Condition and event definitions
04-296r1	Numbering conventions
04-313r1	Mandatory REPORT LUNS Support
05-257r0	Agreed editorial corrections
05-311r0	REMOVE PARTITION Security Controlled by Root Attributes
05-314r1	Possible inconsistencies in integrity check value algorithm field definition/usage
05-316r1	Multi-Object LIST and LIST COLLECTION command enhancements
05-328r1	Four New Multi-Object Commands
06-257r0	Remove PREVENT ALLOW MEDIUM REMOVAL
06-259r2	Making linked commands obsolete

Since the glossary entries described in 04-290r1 (Condition and event definitions) did not exist no new entries were added. The references associated with all the terms identified in 04-290r1 were changed from SAM-3 to SAM-4.

The table showing the general model for the Data-In Buffer and Data-Out Buffer was changed to be more like similar tables in the same level 1 subclause. Specifically, the number of bytes in the integrity check value segment was made variable because the number of integrity check value segment bytes differs between for the Data-In Buffer and Data-Out Buffer.

The attribute length returned for undefined attributes was changed from FFFF FFFFh to FFFFh because the length field in which the value is returned contains only two bytes.

The SCSI documents relationships figure in clause 1 was modified to match the figure in SPC-4. The persistent reservations allowed/conflicts table column headers and table key were updated to match those in SPC-4. The subclause that describes the request-response model was updated to be consistent with SPC-4.

Revision 1 also incorporates changes made by the ANSI Editor that were too editorial to be discussed prior to the development of in revision 0.

The use of italics in normative references was updated to match current practice.

The normative references subclause was updated to contain new ANSI and INCITS contact information as requested by the ANSI Editor.

The list of standards in the SCSI family was removed. This information now appears only in SPC.

2.3 Revision 2 (25 July 2007)

The following approved documents have been incorporated in this revision:

07-257r1	Prohibited needed as a keyword in SPC-4
07-270r2	Several OSD-2 Corrections and Clarifications
07-273r3	Attributes Enhancements
07-274r1	CLEAR command, PUNCH command, & range-based FLUSH
07-275r2	Task Management Function Catchup

During the incorporation of 07-274r1, a few editorial changes were made in the FLUSH, FLUSH COLLECTION, FLUSH OSD, and FLUSH PARTITION commands so that all four have the same wording (e.g., 'the device server

shall ensure are written') as was approved for FLUSH in 07-274r1. Also scope table headers were tweaked so that the FLUSH COLLECTION table did not need to move onto a new page.

Based on editorial comments in 07-105r2, all uses of the verb 'to comprise' were reviewed and ambiguous uses were changed to a variant of the verb 'to consist'.

Based on the new definition of the 'prohibited' keyword in 07-257r1 (all T10 standards use exactly the same keywords), all previous uses of 'prohibited' that did not match the new keyword were revised.

Wherever appropriate 'CHECK CONDITION' was changed to 'CHECK CONDITION status'.

4.6.6.2 (Commands that use collections to affect multiple user objects) was updated to clarify the cases where its statements do not apply to the REMOVE MEMBER OBJECTS command, and that policy tag updates are not always necessary when errors are detected.

In table 51 (Commands for OSD type devices), the footnote mandating support of otherwise optional commands if collections are supported was added to the LIST COLLECTION command.

For the sake of consistency with table 51 (Commands for OSD type devices), all SPC-3 references in the PERFORM SCSI COMMAND definition were changed to SPC-4 references.

In table 84 (LIST COLLECTION command object descriptor format field values), a mistyped 'partition' was changed to 'collection'.

In 4.10.9.2 (Computing updated generation keys and new authentication keys), a reference to the SET MASTER command was corrected to SET MASTER KEY.

The note in 4.12.2 (OSD meta data) was rewritten to clarify the reasons for segmenting the Data-In Buffer and Data-Out Buffer.

In 7.1.2.12 (Root Quotas attributes page), the partition count attribute number in the body text was corrected to match the value in the attributes page definition table.

Several missing and/or incorrect words (e.g., missing 'from') were corrected.

2.4 Revision 3 (22 January 2008)

The following approved documents have been incorporated in this revision:

- 07-301r5 OSD-2 Security Enhancements
- 07-357r0 Correct an OSD-2 defined attributes bug
- 07-378r2 OSD-2 Exceptions Management enhancements
- 08-041r1 Use period as decimal separator in T10 standards

In the capability format definition subclause, the error checking on the SECURITY METHOD field and CAPABILITY FORMAT field was updated to match default security requirements that were changed in r02.

A paragraph that clarifies the applicability of the multi-object rules in 4.6.6.2 was inserted in 4.7.4, (the subclause that describes command function ordering with respect to retrieving and setting attributes).

The table that shows the common CDB format was updated to show where the ALLOCATION LENGTH field, if any, appears.

In the list type values table in clause 7, the description for list type 9h was updated to include all its functions.

Several very minor corrections (e.g., making field names smallcaps) were made but not marked with change bars.

2.5 Revision 4 (24 July 2008)

The following approved documents have been incorporated in this revision:

- 08-179r3 Fixes for five OSD-2 bugs
- 08-181r0 Set Attributes error handling in OSD-2
- 08-182r3 Snapshots and related enhancements
- 08-185r5 CDB Continuations Definition and Usage

The numeric conventions were updated to use the T10 Style Guide text for the use of 'a to z' to represent a range of values. By in large, no other changes were needed in this regard because this standard already contained the notation.

To avoid confusion with defined attributes, all instances of '... defined in x.y.z' were changed to '... described in x.y.z'. These changes were not marked with change bars.

All instances of 'a CHECK CONDITION status' were replaced with 'CHECK CONDITION status'. All instances of 'setting the sense key to' and 'the sense key shall be set to' were replaced with 'with the sense key set to'. All instances of 'the additional sense code shall be set to' were replaced with 'the additional sense key set to'. These changes were not marked with change bars.

Several typographical errors were corrected. These changes were not marked with change bars.

**American National Standards
for Information Systems -**

SCSI Object-Based Storage Device Commands -2 (OSD-2)

Secretariat
InterNational Committee for Information Technology Standards

Approved mm dd yy

American National Standards Institute, Inc.

Abstract

This SCSI command set is designed to provide efficient operation of input/output logical units that manage the allocation, placement, and accessing of variable-size data-storage containers, called objects. Objects are intended to contain operating system and application constructs.

American National Standard

Approval of an American National Standard requires verification by ANSI that the requirements for due process, consensus, and other criteria for approval have been met by the standards developer. Consensus is established when, in the judgment of the ANSI Board of Standards Review, substantial agreement has been reached by directly and materially affected interests. Substantial agreement means much more than a simple majority, but not necessarily unanimity. Consensus requires that all views and objections be considered and that effort be made toward their resolution.

The use of American National Standards is completely voluntary; their existence does not in any respect preclude anyone, whether he or she has approved the standards or not, from manufacturing, marketing, purchasing, or using products, processes, or procedures not confirming to the standards.

The American National Standards Institute does not develop standards and will in no circumstances give interpretation on any American National Standard in the name of the American National Standards Institute. Requests for interpretations should be addressed to the secretariat or sponsor whose name appears on the title page of this standard.

CAUTION NOTICE: This American National Standard may be revised or withdrawn at any time. The procedures of the American National Standards Institute require that action be taken periodically to reaffirm, revise, or withdraw this standard. Purchasers of American National Standards may receive current information on all standards by calling or writing the American National Standards Institute.

CAUTION: The developers of this standard have requested that holders of patents that may be required for the implementation of the standard, disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this standard.

As of the date of publication of this standard, following calls for the identification of patents that may be required for the implementation of the standard, notice of one or more claims has been received.

By publication of this standard, no position is taken with respect to the validity of this claim or of any rights in connection therewith. The known patent holder(s) has (have), however, filed a statement of willingness to grant a license under these rights on reasonable and nondiscriminatory terms and conditions to applicants desiring to obtain such a license. Details may be obtained from the publisher.

No further patent search is conducted by the developer or the publisher in respect to any standard it processes. No representation is made or implied that licenses are not required to avoid infringement in the use of this standard.

Published by
American National Standards Institute
11 West 42nd Street, New York, NY 10036

Copyright 7/24/08 by American National Standards Institute
All rights reserved.

Draft

Printed in the United States of America

Contents

	Page
Foreword	xxii
Introduction	xxiv
1 Scope	1
2 Normative references	3
2.1 Normative references	3
2.2 Approved ISO and ANSI references	3
2.3 Approved FIPS references	3
2.4 Approved IETF References	4
2.5 References under development	4
3 Definitions, symbols, abbreviations, and conventions	5
3.1 Definitions	5
3.2 Acronyms	8
3.3 Keywords	9
3.4 Conventions	10
3.5 Bit and byte ordering	11
3.6 Signed and unsigned integers values	11
3.7 Notation conventions	12
3.7.1 Notation for byte encoded character strings	12
3.7.2 Notation for procedure calls	12
3.8 Data field requirements	13
3.8.1 ASCII data field requirements	13
3.8.2 Data field termination and padding requirements	13
4 SCSI OSD Model	14
4.1 The request-response model	14
4.2 OSD type devices	14
4.3 OSD object abstraction	15
4.4 Elements of the example configuration	16
4.5 Description of the OSD Architecture	17
4.6 Stored data objects	17
4.6.1 Stored data object types	17
4.6.2 Identifying OSD objects	18
4.6.3 Root object	18
4.6.4 Partitions	18
4.6.5 User objects	19
4.6.6 Collections	19
4.6.6.1 Overview	19
4.6.6.2 LINKED collections	20
4.6.6.3 TRACKING collections	20
4.6.6.4 SPONTANEOUS collections	20
4.6.6.5 Well known collections	21
4.6.6.5.1 Overview	21
4.6.6.5.2 The all user objects in partition well known collection	21
4.6.6.5.3 The snapshot/clone tracking well known collection	21
4.6.6.6 Commands that use collections to affect multiple user objects	22
4.7 Data object accessibility	24
4.8 OSD object attributes	25

4.8.1 Overview.....	25
4.8.2 Attribute length when retrieving defined and undefined attributes.....	26
4.8.3 Attribute length when setting defined and undefined attributes.....	27
4.8.4 Command function ordering for commands that get and/or set attributes.....	27
4.8.5 Attributes pages.....	29
4.8.6 Attributes	30
4.8.7 Attributes directories.....	30
4.9 Command atomicity and isolation.....	31
4.9.1 Overview.....	31
4.9.2 Atomicity	31
4.9.3 Isolation	33
4.10 Quotas	33
4.10.1 Introduction.....	33
4.10.2 Quota errors	34
4.10.3 Quota testing	34
4.10.4 Changing quotas	34
4.11 Policy/storage management.....	35
4.11.1 Overview.....	35
4.11.2 Capabilities	35
4.11.2.1 Introduction.....	35
4.11.2.2 Capability format.....	36
4.11.2.2.1 Introduction.....	36
4.11.2.2.2 USER capability object descriptor	41
4.11.2.2.3 PAR capability object descriptor.....	43
4.11.2.2.4 COL capability object descriptor.....	44
4.11.2.2.5 Command termination due to errors detected in a capability	45
4.11.2.3 Capabilities and commands allowed	45
4.11.3 OBSD storage damage detection, repair, and undesirable utilization prevention	53
4.11.3.1 Normal usage storage damage detection and repair	53
4.11.3.2 Policy access tags	54
4.11.3.3 Storage damage detection and repair after a reset.....	54
4.12 Security.....	56
4.12.1 Basic security model.....	56
4.12.2 Trust assumptions	57
4.12.3 Preparing credentials.....	58
4.12.3.1 Introduction.....	58
4.12.3.2 Capability setup steps for credential preparation	59
4.12.4 Security methods.....	60
4.12.4.1 Introduction.....	60
4.12.4.2 The NOSEC security method	61
4.12.4.3 The CAPKEY security method	62
4.12.4.4 The CMDRSP security method	62
4.12.4.5 The ALLDATA security method.....	63
4.12.5 Credentials	67
4.12.5.1 Credential format.....	67
4.12.5.2 Capability keys	68
4.12.6 OSD device server security algorithms	68
4.12.6.1 Credential validation	68
4.12.6.1.1 Introduction.....	68
4.12.6.1.2 Validating the solo portion of a credential.....	69
4.12.6.1.3 Validating the extension portion of a credential.....	69
4.12.6.1.4 Validating a CDB continuation segment using the solo portion of a credential	70
4.12.6.2 Reconstructing the credential	70
4.12.6.2.1 Reconstructing the solo portion of a credential	70

4.12.6.2.2 Reconstructing the extension portion of a credential	71
4.12.6.3 Computing the integrity check values for a credential	71
4.12.6.3.1 Computing the solo integrity check value for a credential	71
4.12.6.3.2 Computing the extended integrity check value for a credential	72
4.12.6.4 Invalidating credentials	72
4.12.7 Request nonces	73
4.12.7.1 Request nonce format	73
4.12.7.2 Device server validation of request nonces	73
4.12.7.3 Lists of previously used request nonces	74
4.12.7.3.1 Introduction	74
4.12.7.3.2 Freezing capability audit fields	74
4.12.7.3.3 Freezing working keys	75
4.12.8 Integrity check values	75
4.12.9 Secret keys	76
4.12.9.1 Introduction	76
4.12.9.2 Computing updated generation keys and new authentication keys	77
4.12.10 OSD security interactions with SPC-3 commands and SAM-3 task management functions	78
4.13 Object duplication	78
4.13.1 Overview	78
4.13.2 Snapshot partitions and clone partitions	78
4.13.2.1 Overview	78
4.13.2.2 Snapshot history chains	79
4.13.2.3 Clone chains	81
4.13.2.4 Snapshots of clones	84
4.13.2.5 Detaching a clone partition	85
4.13.3 Object duplication methods	86
4.13.4 Object duplication state management	87
4.13.4.1 Overview	87
4.13.4.2 Time of duplication source object management	87
4.13.4.3 Source object freeze duplication management	87
4.13.5 Object duplication space accounting	87
4.14 Data persistence model	88
4.15 Data-In and Data-Out Buffer model	88
4.15.1 Bidirectional data transfers	88
4.15.2 OSD meta data	88
4.15.3 OSD Data-In Buffer format	90
4.15.4 OSD Data-Out Buffer format	91
4.15.5 Data-In and Data-Out buffer offsets	92
4.16 Error reporting	93
4.16.1 Introduction	93
4.16.2 OSD-specific sense data descriptors	94
4.16.2.1 OSD error identification sense data descriptor	94
4.16.2.2 OSD response integrity check value sense data descriptor	96
4.16.2.3 OSD attribute identification sense data descriptor	96
4.16.3 Auto contingent allegiance	97
4.17 Reservations	97
5 Common Formats	100
5.1 OSD CDB format	100
5.2 Fields commonly used in OSD commands	101
5.2.1 Overview	101
5.2.2 Allocation length	102
5.2.3 Caching control bits	102
5.2.4 Capability	102

5.2.5 CDB continuation length	102
5.2.6 Get and set attributes parameters	103
5.2.6.1 Get and set attributes CDB format selection	103
5.2.6.2 Set one attribute value using CDB fields	103
5.2.6.3 Get an attributes page and set an attribute value	104
5.2.6.4 Get and set attributes lists	106
5.2.7 Immediate bit for TRACKING collections	107
5.2.8 Isolation	108
5.2.9 Length	108
5.2.10 Partition_ID	108
5.2.11 Security parameters	109
5.2.12 Starting byte address	110
5.2.13 Timestamps control	110
5.2.14 User_Object_ID	110
5.3 CDB continuation segment format	111
5.4 CDB continuation descriptors	112
5.4.1 Overview	112
5.4.2 Scatter/gather list	114
5.4.3 Query list	116
5.4.4 User object	118
5.4.5 Copy user object source	119
5.4.6 Extension capabilities CDB continuation descriptors	121
6 Commands for OSD type devices	122
6.1 Summary of commands for OSD type devices	122
6.2 APPEND	125
6.3 CLEAR	127
6.4 COPY USER OBJECTS	129
6.5 CREATE	131
6.6 CREATE AND WRITE	134
6.7 CREATE CLONE	136
6.7.1 Introduction	136
6.7.2 Processing before the IMMED_TR bit takes effect	138
6.7.3 Processing after the IMMED_TR bit takes effect, if any	140
6.7.4 Command completion	140
6.8 CREATE COLLECTION	141
6.9 CREATE PARTITION	143
6.10 CREATE SNAPSHOT	145
6.10.1 Introduction	145
6.10.2 Processing before the IMMED_TR bit takes effect	147
6.10.3 Processing after the IMMED_TR bit takes effect, if any	149
6.10.4 Command completion	149
6.11 CREATE USER TRACKING COLLECTION	150
6.12 DETACH CLONE	153
6.13 FLUSH	156
6.14 FLUSH COLLECTION	158
6.15 FLUSH OSD	160
6.16 FLUSH PARTITION	162
6.17 FORMAT OSD	164
6.18 GET ATTRIBUTES	166
6.19 GET MEMBER ATTRIBUTES	167
6.20 LIST	169
6.20.1 Introduction	169
6.20.2 LIST command parameter data	172

6.20.3 LIST command and LIST COLLECTION command object descriptor formats	173
6.20.3.1 Partition_ID only object descriptor format.....	173
6.20.3.2 Collection_Object_ID only object descriptor format.....	174
6.20.3.3 User_Object_ID only object descriptor format.....	174
6.21 LIST COLLECTION	175
6.22 OBJECT STRUCTURE CHECK.....	180
6.22.1 Introduction.....	180
6.22.2 Structure checking for the root object and all partitions.....	181
6.22.3 Structure checking for a specific partition.....	182
6.23 PERFORM SCSI COMMAND	183
6.24 PERFORM TASK MANAGEMENT FUNCTION.....	186
6.25 PUNCH.....	189
6.26 QUERY.....	191
6.26.1 Introduction.....	191
6.26.2 Matches list parameter data format	193
6.27 READ.....	194
6.28 READ MAP.....	196
6.28.1 Introduction.....	196
6.28.2 READ MAP command and READ MAPS AND COMPARE command parameter data.....	198
6.29 READ MAPS AND COMPARE.....	201
6.30 REFRESH SNAPSHOT OR CLONE.....	203
6.30.1 Introduction.....	203
6.30.2 Processing before the IMMED_TR bit takes effect.....	205
6.30.3 Processing after the IMMED_TR bit takes effect, if any	206
6.30.4 Command completion.....	207
6.30.5 Linking a snapshot as the most recent entry in the history chain	208
6.30.6 Unlinking a snapshot from the history chain.....	208
6.31 REMOVE	209
6.32 REMOVE COLLECTION	210
6.33 REMOVE MEMBER OBJECTS.....	212
6.34 REMOVE PARTITION.....	213
6.35 RESTORE PARTITION FROM SNAPSHOT.....	215
6.35.1 Introduction.....	215
6.35.2 Processing before the IMMED_TR bit takes effect.....	216
6.35.3 Processing after the IMMED_TR bit takes effect, if any	217
6.35.4 Command completion.....	217
6.36 SET ATTRIBUTES	219
6.37 SET KEY	220
6.38 SET MASTER KEY	222
6.38.1 Introduction.....	222
6.38.2 Seed exchange.....	223
6.38.3 Change master key	224
6.39 SET MEMBER ATTRIBUTES	227
6.40 WRITE	229
7 Parameters for OSD type devices.....	231
7.1 Attributes parameters	231
7.1.1 Attributes parameter formats.....	231
7.1.2 Reporting illegal attempts to set attribute values.....	231
7.1.3 OSD attributes pages	232
7.1.3.1 Attributes pages overview	232
7.1.3.2 Attribute number 0h in all attributes pages.....	233
7.1.3.3 Attribute number 0h for unidentified attributes pages.....	234
7.1.3.4 Root Directory attributes page.....	235

7.1.3.5 Partition Directory attributes page	236
7.1.3.6 Collection Directory attributes page.....	237
7.1.3.7 User Object Directory attributes page	238
7.1.3.8 Root Information attributes page	239
7.1.3.9 Partition Information attributes page.....	247
7.1.3.10 Collection Information attributes page	250
7.1.3.11 User Object Information attributes page.....	252
7.1.3.12 Root Quotas attributes page.....	254
7.1.3.13 Partition Quotas attributes page	256
7.1.3.14 User Object Quotas attributes page	257
7.1.3.15 Root Timestamps attributes page.....	259
7.1.3.16 Partition Timestamps attributes page	261
7.1.3.17 Collection Timestamps attributes page	263
7.1.3.18 User Object Timestamps attributes page	265
7.1.3.19 Attributes Access attributes page	266
7.1.3.20 Command Tracking attributes page	268
7.1.3.21 Collections attributes page	270
7.1.3.22 Root Policy/Security attributes page.....	272
7.1.3.23 Partition Policy/Security attributes page	278
7.1.3.24 Collection Policy/Security attributes page	282
7.1.3.25 User Object Policy/Security attributes page	283
7.1.3.26 Root Error Recovery attributes page	284
7.1.3.27 Partition Error Recovery attributes page	288
7.1.3.28 Collection Error Recovery attributes page	291
7.1.3.29 User Object Error Recovery attributes page.....	292
7.1.3.30 Snapshots Information attributes page.....	295
7.1.3.31 Current Command attributes page	298
7.1.3.32 Null attributes page.....	300
7.1.4 OSD attributes lists.....	300
7.1.4.1 Attributes lists overview	300
7.1.4.2 List entry format for retrieving attributes for a specified OSD object	302
7.1.4.3 List entry format for retrieved attributes and for setting attributes for the specified OSD object	303
7.1.4.4 Multi-object retrieved attributes format	304
7.2 Diagnostic parameters.....	305
7.3 Log parameters	305
7.4 Mode parameters	305
7.5 Vital product data parameters	306
7.5.1 Overview.....	306
7.5.2 OSD Information VPD page	306
7.5.2.1 Overview.....	306
7.5.2.2 OSD logical unit security methods information descriptor	307
7.5.3 Security Token VPD page	308
Annex A (Normative) Attributes page numbers assigned by other standards	309
A.1 Attributes page numbers assigned by other standards.....	309
Annex B (Informative) Numeric order codes	310
B.1 Service action codes	310
Annex C (Informative) Attributes defined by this standard.....	312
C.1 Attributes list	312
Annex D (Informative) Examples of OSD Operation.....	318
D.1 Preparing a device for OSD operation	318

D.2 Example of accessing data on an OSD	319
---	-----

Tables

	Page
1 Numbering conventions examples	10
2 OSD model objects	17
3 Partition_ID and User_Object_ID value assignments	18
4 TRACKING collection operational characteristics	20
5 Well known collections	21
6 Attributes retrieval and setting requirements for multi-object commands	24
7 DATA PROTECT INFORMATION field contents	25
8 Setting defined attributes and undefined attributes	27
9 Attributes page numbers	29
10 Attributes page number sets	29
11 Attributes directory pages	30
12 Atomicity attributes	32
13 Examples of atomicity attributes effects	33
14 Capability format	36
15 Capability format values	36
16 Created time for OSD objects by type	38
17 Object type values	38
18 Permissions bit mask format	38
19 Object descriptor types	40
20 ALLOWED ATTRIBUTES ACCESS field	40
21 User object descriptor format	41
22 Policy access tag usage for OSD object types and commands	41
23 Partition descriptor format	43
24 Collection descriptor format	44
25 Commands allowed by specific capability field values	45
26 Attribute retrieving and setting function allowed by specific capability field values	50
27 Policy access tag format	54
28 Security manager communications trust requirements	58
29 OSD security methods	60
30 Security methods and threats thwarted	61
31 Data-out integrity information format	64
32 Data-in integrity information format	66
33 Credential format	67
34 Request nonce format	73
35 OSD secret key hierarchy	76
36 Comparison of snapshots and clones	79
37 Snapshot history chain attributes for one CREATE SNAPSHOT command	80
38 Snapshot history chain attributes for two CREATE SNAPSHOT commands	81
39 Snapshot/clone chain attributes for one CREATE CLONE command	82
40 Snapshot/clone chain attributes for two CREATE CLONE commands	83
41 Snapshot/clone chain attributes for a CREATE SNAPSHOT command on a clone partition	84
42 Snapshot/clone chain attributes after a DETACH CLONE command	85
43 Object duplication methods	86
44 Time of duplication source object management	87
45 OSD Data-In Buffer and Data-Out Buffer model	89
46 OSD Data-In Buffer format	90
47 Summary of OSD Data-In Buffer offsets	90
48 OSD Data-Out Buffer format	91
49 Summary of OSD Data-Out Buffer offsets	92
50 CDB Data-In Buffer and Data-Out Buffer offset field format	92
51 OSD object identification sense data descriptor format	94
52 Command functions bits	95

53 Command functions indicated by the command functions bits	95
54 Command functions bits combinations	95
55 OSD response integrity check value sense data descriptor format	96
56 OSD attribute identification sense data descriptor format.....	96
57 Sense data attribute descriptor format.....	97
58 OSD commands that are allowed in the presence of various reservations.....	98
59 Basic OSD CDB.....	100
60 OSD service action specific fields	101
61 Get and set attributes CDB format code values.....	103
62 Set one attribute using CDB parameters format	103
63 Page oriented get and set attributes CDB parameters format	104
64 List oriented get and set attributes CDB parameters format.....	106
65 ISOLATION field	108
66 Security parameters format.....	109
67 Timestamps control values	110
68 CDB continuation segment format.....	111
69 CDB CONTINUATION FORMAT field	111
70 CDB continuation descriptor format.....	112
71 CDB CONTINUATION DESCRIPTOR TYPE field	113
72 Scatter/gather list CDB continuation descriptor format	114
73 Scatter/gather list entry format.....	115
74 Query list CDB continuation descriptor format.....	116
75 QUERY TYPE field values	117
76 Query criteria entry format	117
77 User object CDB continuation descriptor format.....	118
78 Copy user object source CDB continuation descriptor format	119
79 Extension capabilities CDB continuation descriptor format.....	121
80 Commands for OSD type devices.....	122
81 APPEND command	125
82 CLEAR command	127
83 COPY USER OBJECTS command	129
84 CREATE command.....	131
85 CREATE AND WRITE command.....	134
86 CREATE CLONE command	136
87 CREATE COLLECTION command.....	141
88 CREATE PARTITION command	143
89 CREATE SNAPSHOT command.....	145
90 CREATE USER TRACKING COLLECTION command	150
91 DETACH CLONE command	153
92 FLUSH command	156
93 User object flush scope values	157
94 FLUSH COLLECTION command	158
95 Collection flush scope values.....	159
96 FLUSH OSD command.....	160
97 Root object flush scope values	161
98 FLUSH PARTITION command	162
99 Partition flush scope values	163
100 FORMAT OSD command	164
101 GET ATTRIBUTES command	166
102 GET MEMBER ATTRIBUTES command.....	167
103 LIST command.....	169
104 Specifying objects and attributes to be listed	170
105 LIST command sort order values.....	170
106 Attributes processing requirements for LIST commands with the LIST_ATTR bit set to one.....	171

107 LIST command parameter data	172
108 LIST command OBJECT DESCRIPTOR FORMAT field values	173
109 Partition_ID only object descriptor format	173
110 Collection_Object_ID only object descriptor format	174
111 User_Object_ID only object descriptor format	174
112 LIST COLLECTION command	175
113 Specifying collections and attributes to be listed	176
114 Attributes processing requirements for LIST COLLECTION commands with the LIST_ATTR bit set to one....	177
115 LIST COLLECTION command parameter data	178
116 LIST COLLECTION command OBJECT DESCRIPTOR FORMAT field values	179
117 OBJECT STRUCTURE CHECK command	180
118 PERFORM SCSI COMMAND command	183
119 Request CDBs allowed in the PERFORM SCSI COMMAND	184
120 PERFORM TASK MANAGEMENT FUNCTION command	186
121 TASK MANAGEMENT FUNCTION field	187
122 PERFORM TASK MANAGEMENT FUNCTION parameter data format	188
123 SERVICE RESPONSE field	188
124 PUNCH command	189
125 QUERY command	191
126 Matches list parameter data format	193
127 READ command	194
128 READ MAP command	196
129 REQUESTED MAP TYPE field	197
130 READ MAP command and READ MAPS AND COMPARE command parameter data	198
131 Map descriptor format	198
132 MAP DESCRIPTOR INDEX field	199
133 MAP DESCRIPTOR TYPE field	199
134 READ MAPS AND COMPARE command	201
135 COMPARISON SCOPE field	202
136 REFERESH SNAPSHOT OR CLONE command	203
137 Snapshot backward and forward attribute values set to add the most recent entry	208
138 Snapshot backward and forward attribute values set to unlink an entry	208
139 REMOVE command	209
140 REMOVE COLLECTION command	210
141 REMOVE MEMBER OBJECTS command	212
142 REMOVE PARTITION command	213
143 REMOVE SCOPE field	214
144 RESTORE PARTITION FROM SNAPSHOT command	215
145 SET ATTRIBUTES command	219
146 SET KEY command	220
147 Key to set code values	221
148 SET MASTER KEY command	222
149 Diffie-Hellman exchange step values	223
150 Seed exchange device server DH_data format	224
151 Change master key DH_data format	225
152 SET MEMBER ATTRIBUTES command	227
153 WRITE command	229
154 Attributes pages defined by this standard	232
155 Attribute number 0h format for all attributes pages	233
156 Example Root Directory attributes page contents	235
157 Example Partition Directory attributes page contents	236
158 Example Collection Directory attributes page contents	237
159 Example User Object Directory attributes page contents	238
160 Root Information attributes page contents	239

161 Object accessibility attribute values	241
162 Default isolation method attribute values	242
163 Supported isolation methods attribute contents	242
164 Supported object duplication method attributes contents	243
165 Supported time of duplication method attributes contents	244
166 Support for duplicated object freezing attribute contents	245
167 Support for snapshot refreshing attribute values	245
168 Supported CDB continuation descriptor type attribute contents	246
169 Partition Information attributes page contents	247
170 Collection Information attributes page contents	250
171 Collection type codes	251
172 User Object Information attributes page contents	252
173 Root Quotas attributes page contents	254
174 Root Quotas attributes page format	255
175 Partition Quotas attributes page contents	256
176 Partition Quotas attributes page format	257
177 User Object Quotas attributes page contents	257
178 User Object Quotas attributes page format	258
179 Root Timestamps attributes page contents	259
180 Timestamp bypass attribute values	259
181 Root Timestamps attributes page format	260
182 Partition Timestamps attributes page contents	261
183 Partition Timestamps attributes page format	262
184 Collection Timestamps attributes page contents	263
185 Collection Timestamps attributes page format	264
186 User Object Timestamps attributes page contents	265
187 User Object Timestamps attributes page format	266
188 Attributes Access attributes page contents	266
189 Allowed attributes access attribute format	267
190 Attributes access descriptor	267
191 Command Tracking attributes page contents	268
192 Active command status attribute values	269
193 Ended command status attribute values	269
194 Collections attributes page contents	270
195 Collections attributes page format	271
196 Root Policy/Security attributes page contents	272
197 Supported security methods attribute format	273
198 Supported integrity check value algorithm codes	274
199 Root Policy/Security attributes page format	276
200 Partition Policy/Security attributes page contents	278
201 Frozen working key bit mask attribute format	279
202 Partition Policy/Security attributes page format	281
203 Collection Policy/Security attributes page contents	282
204 Collection Policy/Security attributes page format	283
205 User Object Policy/Security attributes page contents	283
206 User Object Policy/Security attributes page format	284
207 Root Error Recovery attributes page contents	284
208 Root damage summary attribute value	285
209 Contained objects damage summary root attribute value	285
210 Root Error Recovery attributes page format	287
211 Partition Error Recovery attributes page contents	288
212 Partition damage summary attribute value	288
213 Contained objects damage summary partition attribute value	289
214 Partition Error Recovery attributes page format	290

215 Collection Error Recovery attributes page contents.....	291
216 Collection damage summary attribute value	291
217 Collection Error Recovery attributes page format	292
218 User Object Error Recovery attributes page contents.....	292
219 User object damage summary attribute value	293
220 User Object Error Recovery attributes page format.....	294
221 Snapshots Information attributes page contents.....	295
222 Partition type attribute values.....	296
223 Current Command attributes page contents	298
224 Current Command attributes page format	299
225 Null attributes page format	300
226 Attributes list format	300
227 List type values	301
228 List entry format for retrieving attributes for a specified OSD object.....	302
229 List entry format for retrieved attributes and for setting attributes for the specified OSD object.....	303
230 Multi-object retrieved attributes list format	304
231 OSD specific VPD page codes	306
232 OSD Information VPD page format.....	306
233 OSD information descriptor format	307
234 OSD information descriptor type values	307
235 OSD logical unit security methods information descriptor format	307
236 Security Token VPD page.....	308
A.1 Attributes page numbers assigned by other standards	309
B.1 Numerical order OSD service action codes.....	310
C.1 Numerical order attributes defined by this standard.....	312
D.1 OSD initialization sequence	318
D.2 OSD command sequence for creating a file.....	319
D.3 OSD command sequence using CREATE AND WRITE	319

Figures

	Page
1 SCSI document relationships.....	1
2 Comparison of traditional and OSD storage models.....	15
3 Example OSD Configuration.....	16
4 OSD security model transactions.....	56
5 Snapshot history chain after first CREATE SNAPSHOT command.....	80
6 Snapshot history chain after second CREATE SNAPSHOT command.....	81
7 Snapshot/clone chains after a first CREATE CLONE command.....	82
8 Snapshot/clone chains after a second CREATE CLONE command.....	83
9 Snapshot/clone chains after a CREATE SNAPSHOT command on a clone partition.....	84
10 Snapshot/clone chains after a DETACH CLONE command.....	85

Foreword

This foreword is not part of American National Standard INCITS.***:200x.

This SCSI command set is designed to provide efficient operation of input/output logical units that manage the allocation, placement, and accessing of variable-size data-storage containers, called objects. Objects are intended to contain operating system and application constructs.

This SCSI command set provides multiple operating systems concurrent control over one or more logical units. However, the multiple operating systems are assumed to properly coordinate their actions to prevent data corruption. This SCSI standard provides commands that assist with coordination between multiple operating systems. However, details of the coordination are beyond the scope of this SCSI command set.

This standard defines a logical unit model for Object-Based Storage Device logical units. Also defined are SCSI commands that apply to Object-Based Storage Device logical units.

Objects designate entities in which computer systems store data. The purpose of this abstraction is to assign to the storage device the responsibility for managing where data is located on the device.

This standard was developed by T10 in cooperation with industry groups during 1999 through 2004.

With any technical document there may arise questions of interpretation as new products are implemented. INCITS has established procedures to issue technical opinions concerning the standards developed by INCITS. These procedures may result in SCSI Technical Information Bulletins being published by INCITS.

These Bulletins, while reflecting the opinion of the Technical Committee that developed the standard, are intended solely as supplementary information to other users of the standard. This standard, ANSI INCITS.***:200x, as approved through the publication and voting procedures of the American National Standards Institute, is not altered by these bulletins. Any subsequent revision to this standard may or may not reflect the contents of these Technical Information Bulletins.

Current INCITS practice is to make Technical Information Bulletins available through:

INCITS Online Store	http://www.techstreet.com/incits.html
managed by Techstreet	Telephone: 1-734-302-7801 or
1327 Jones Drive	1-800-699-9277
Ann Arbor, MI 48105	Facsimile: 1-734-302-7811

or

Global Engineering	http://global.ihs.com/
15 Inverness Way East	Telephone: 1-303-792-2181 or
Englewood, CO 80112-5704	1-800-854-7179
	Facsimile: 1-303-792-2192

Requests for interpretation, suggestions for improvement and addenda, or defect reports are welcome. They should be sent to the INCITS Secretariat, InterNational Committee for Information Technology Standards, Information Technology Institute, 1250 Eye Street, NW, Suite 200, Washington, DC 20005-3922.

This standard was processed and approved for submittal to ANSI by the InterNational Committee for Information Technology Standards (INCITS). Committee approval of the standard does not necessarily imply that all committee members voted for approval. At the time of it approved this standard, INCITS had the following members:

<<Insert INCITS member list>>

Technical Committee T10 on SCSI Storage Interfaces, which developed and reviewed this standard, had the following members:

John B. Lohmeyer, Chair
Mark Evans, Vice-Chair
Ralph O. Weber, Secretary

The T10 Technical Committee expresses its appreciation to the Storage Networking Industry Association (SNIA, see <http://www.snia.org>) Object-based Storage (ObS) Technical Working Group (TWG) for their contributions to this standard. The SNIA ObS TWG members were:

Mr. Julian Satran, IBM Co-Chair
Dr. Sami Iren, Seagate Technology, Co-Chair

Introduction

The SCSI Object-Based Storage Device Commands -2 (OSD-2) standard is divided into the following clauses and annexes:

- Clause 1 is the scope.
- Clause 2 enumerates the normative references that apply to this standard.
- Clause 3 describes the definitions, symbols, and abbreviations used in this standard.
- Clause 4 describes the model for an OSD device and the conceptual relationship between this document and the SCSI Architecture Model.
- Clause 5 describes the CDB formats used throughout this standard.
- Clause 6 describes commands that may be implemented by a SCSI device that conforms to this standard.
- Clause 7 defines the parameter data formats that may be implemented by a SCSI device that conforms to this standard.
- Annex A lists attributes page numbers assigned by other standards.
- Annex B lists OSD service actions in numerical order.
- Annex C lists attributes defined by this standard in numerical order.
- Annex D gives examples of OSD usage.

American National Standard for Information Systems -
Information Technology -
SCSI Object-Based Storage Device Commands -2 (OSD-2)

1 Scope

This standard defines the command set extensions to control operation of Object-Based Storage devices. The clause(s) of this standard pertaining to the SCSI Object-Based Storage Device class, implemented in conjunction with the applicable clauses of the ISO/IEC 14776-453 SCSI Primary Commands -3 (SPC-3), specify the standard command set for SCSI Object-Based Storage devices.

The objective of this standard is to provide the following:

- a) Permit an application client to communicate with a logical unit that declares itself to be an Object-Based Storage device in the PERIPHERAL DEVICE TYPE field of the INQUIRY command response data over a SCSI service delivery subsystem;
- b) Enable construction of a shared storage processor cluster with equipment and software from many different vendors;
- c) Define commands unique to the type of SCSI Object-Based Storage devices;
- d) Define commands to manage the operation of SCSI Object-Based Storage devices.

The set of SCSI standards specifies the interfaces, functions, and operations necessary to ensure interoperability between conforming SCSI implementations. This standard is a functional description. Conforming implementations may employ any design technique that does not violate interoperability.

Figure 1 shows the relationship of this standard to the other standards and related projects in the SCSI family of standards as of the publication of this standard.

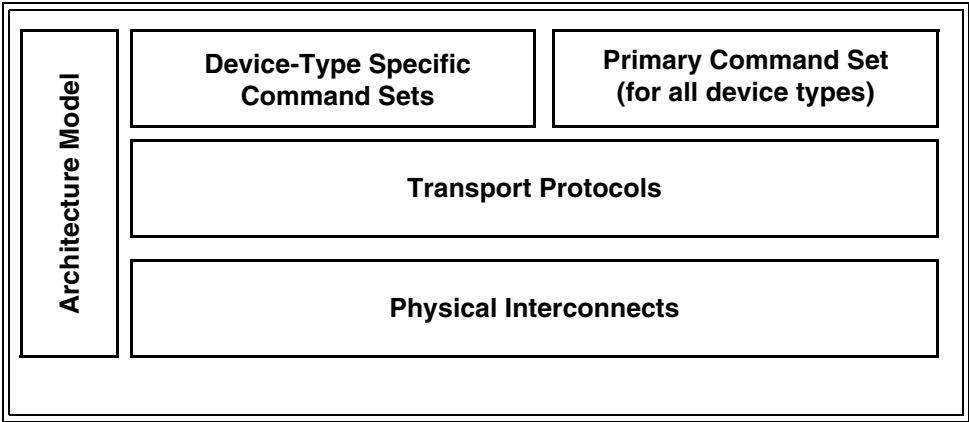


Figure 1 — SCSI document relationships

Figure 1 is intended to show the general relationship of the documents to one another. Figure 1 is not intended to imply a relationship such as a hierarchy, protocol stack, or system architecture. It indicates the applicability of a standard to the implementation of a given transport.

The term SCSI is used to refer to the family of standards described in this clause.

The following features from previous standards have been made obsolete in this standard:

- a) The ROOT bit in the LIST command parameter data;
- b) The COLTN bit in the LIST COLLECTION command parameter data;
- c) All discussion of linked commands;
- d) Since the CDB format was changed globally, all service action coded value assignments from OSD were made obsolete in this standard and new coded values were assigned.

2 Normative references

2.1 Normative references

The standards identified in this subclause contain provisions that, by reference in the text, constitute provisions of this standard. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent editions of the standards listed in this subclause.

2.2 Approved ISO and ANSI references

Copies of the following documents may be obtained from ANSI, an ISO member organization:

- a) Approved ANSI standards;
- b) approved international and regional standards (ISO and IEC); and
- c) approved foreign standards (including JIS and DIN).

For further information, contact the ANSI Customer Service Department:

Phone +1 212-642-4900
Fax: +1 212-302-1286
Web: <http://www.ansi.org>
E-mail: ansionline@ansi.org

or the InterNational Committee for Information Technology Standards (INCITS):

Phone +1 202-626-5738
Web: <http://www.incits.org>
E-mail: incits@itic.org

ISO/IEC 14776-413, *SCSI Architecture Model - 3 (SAM-3)* [ANSI INCITS 402-2005]

ISO/IEC 14776-453, *SCSI Primary Commands - 3 (SPC-3)* [ANSI INCITS 408-2005]

ANSI INCITS 400-2004, *Object-based Storage Device Commands (OSD)*

2.3 Approved FIPS references

Copies of Federal Information Processing Standards (FIPS) document may be obtained via the World Wide Web site (<http://www.itl.nist.gov/fipspubs/>). In the event that FIPS World Wide Web site is no longer active, access may be possible via the Information Technology Laboratory World Wide Web site (<http://www.itl.nist.gov/>) or the National Institute of Standards and Technology site (<http://www.nist.gov/>).

FIPS 180-1 (1995), *Secure Hash Standard (i.e., SHA1)*

FIPS 198 (2002), *The Keyed-Hash Message Authentication Code (HMAC)*

2.4 Approved IETF References

Copies of the following approved IETF standards may be obtained through the Internet Engineering Task Force (IETF) at <http://www.ietf.org>.

RFC 1750, *Randomness Recommendations for Security*

RFC 2401, *Security Architecture for the Internet Protocol*

RFC 2409, *The Internet Key Exchange*

RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange*

2.5 References under development

At the time of publication, the following referenced standards were still under development by T10 (<http://www.t10.org>). For information on the current status of the document, or regarding availability, contact the T10 Technical Committee or INCITS (<http://www.incits.org>).

ISO/IEC 14776-414, *SCSI Architecture Model - 4 (SAM-4)* [T10/1683-D]

ISO/IEC 14776-454, *SCSI Primary Commands - 4 (SPC-4)* [T10/1731-D]

3 Definitions, symbols, abbreviations, and conventions

3.1 Definitions

3.1.1 additional sense code: A combination of the ADDITIONAL SENSE CODE field and ADDITIONAL SENSE CODE QUALIFIER field in the sense data (see 3.1.44).

3.1.2 application client: An object that is the source of SCSI commands. See SAM-3.

3.1.3 attributes: Data, sometimes called meta data, that is associated with an OSD object (see 3.1.29) that is not accessible via read or write command functions (see 3.1.10). See 4.8.

3.1.4 capability: The fields in a CDB or CDB continuation segment (see 5.3) that specify what command functions (see 3.1.10) the command may request (e.g., what OSD object (see 3.1.29) may be accessed). The contents of capabilities may be managed for application clients by a policy/storage manager (see 3.1.34) and secured in credentials (see 3.1.11) by a security manager (see 3.1.41). See 4.11.2.

3.1.5 capability key: An integrity check value (see 3.1.19) computed for one or more capabilities and sent to an application client in a credential (see 3.1.11) that is used by the application client to compute integrity check values for an OSD command. See 4.12.5.2.

3.1.6 collection: An OSD object (see 3.1.29) in which references to one or more user objects from a single partition (see 3.1.31) may be collected. See 4.6.6.

3.1.7 Collection_Object_ID: The identifier for one collection (see 3.1.6).

3.1.8 command: A request describing one or more command functions (see 3.1.10) to be performed by a device server. See SAM-3.

3.1.9 command descriptor block (CDB): The structure used to communicate commands from an application client to a device server. See SPC-3.

3.1.10 command function: One unit of work within a single command (see 3.1.8). This standard extends the SAM-3 definition of command to allow multiple command functions to be requested by a single command.

3.1.11 credential: A data structure that is prepared by the security manager (see 3.1.41) and protected by an integrity check value (see 3.1.19) that is sent to an application client in order to grant defined access to an OSD logical unit for specific command functions (see 3.1.10) performed on specific OSD objects. The credential includes a capability (see 3.1.4) that is prepared by the policy/storage manager (see 3.1.34) that the application client copies to each CDB that requests the specified command functions. See 4.12.5.

3.1.12 Data-In Buffer: The buffer identified by the application client to receive data from the device server during the processing of a command. See SAM-3.

3.1.13 Data-Out Buffer: The buffer identified by the application client to supply data that is sent from the application client to the device server during the processing of a command. See SAM-3.

3.1.14 defined attribute: An attribute (see 3.1.3) for which the OSD logical unit (see 3.1.28) is storing a value of length greater than zero (i.e., an attribute that is known to the OSD logical unit).

3.1.15 device server: An object within a logical unit that processes SCSI tasks according to the rules of task management. See SAM-3.

3.1.16 extension capability: A capability (see 3.1.4) that is not the first capability returned in a credential (see 3.1.11) and is placed in the CDB continuation segment (see 5.3). A credential binds one or more extension capabilities to one solo capability (see 3.1.46).

3.1.17 field: A group of one or more contiguous bits, a part of a larger structure such as a CDB (see 3.1.9) or sense data (see 3.1.44).

3.1.18 I_T nexus: A nexus between a SCSI initiator port and a SCSI target port. See SAM-3.

3.1.19 integrity check value: A value computed using a security algorithm (e.g., HMAC-SHA1), a secret key (see 3.1.40), and an array of bytes. See 4.12.8.

3.1.20 left-aligned: A type of field containing ASCII data in which unused bytes are placed at the end of the field (i.e., highest offset). See 3.8.1.

3.1.21 logical unit: An externally addressable entity within a SCSI device that implements a SCSI device model and contains a device server. See SAM-3.

3.1.22 meta data: Information associated with an object that is not user data (e.g., attributes (see 3.1.3)).

3.1.23 nexus: A relationship between two SCSI devices, and the SCSI initiator port and SCSI target port objects within those SCSI devices. See SAM-3.

3.1.24 nonce: A value that is used one and only one time and thus uniquely identifies a single instance of something (e.g., an individual OSD command, or one credential) transacted between an application client, device server, and security manager.

3.1.25 null-padded: A type of field in which unused bytes are filled with ASCII null (00h) characters. See 3.8.2.

3.1.26 object: 1: An ordered set of bytes within an object-based storage device that is associated with a unique identifier. Data in the object is referenced by the identifier and offset information within the object. Objects are allocated and placed on the media by the OSD logical unit. 2: When used in relationship to SAM-3, a SCSI architecture model object. See SAM-3.

3.1.27 object-based storage device (OBSD): A SCSI device that implements this standard in which data is organized and accessed as objects.

3.1.28 OSD logical unit: A logical unit within an OBSD (see 3.1.27).

3.1.29 OSD object: A root object (see 3.1.36), a partition (see 3.1.31), a collection (see 3.1.6), or user object (see 3.1.53).

3.1.30 page: A regular parameter structure or format used by several commands. These pages are identified with a value known as a page code or page number.

3.1.31 partition: An OSD object (see 3.1.29) used for creating distinct management domains (e.g., for naming, security, quota management). See 4.6.3.

3.1.32 Partition_ID: The identifier for one partition (see 3.1.31).

3.1.33 partition zero: The partition with the Partition_ID (see 3.1.32) zero. The partition that represents the root object (see 3.1.36).

3.1.34 policy/storage manager: The component of an OSD configuration (see 4.4) that manages prevention of unsafe or temporarily undesirable utilization of OBSD (see 3.1.27) storage, coordinates access policies, and prepares capabilities (see 3.1.4) that specify what command functions (see 3.1.10) the command may request (e.g., what OSD object (see 3.1.29) may be accessed). See 4.11.

3.1.35 request nonce: A nonce (see 3.1.24) having the format used by OSD command requests and responses. See 4.12.7.

3.1.36 root object: An OSD object (see 3.1.29) that is always present whose attributes contain global characteristics for the OSD logical unit. Each OSD logical unit has one and only one root object. See 4.6.3.

3.1.37 SCSI device: A device that contains one or more SCSI ports that are connected to a service delivery subsystem and supports a SCSI application protocol. See SAM-3.

3.1.38 SCSI initiator port: A SCSI initiator device object that acts as the connection between application clients and the service delivery subsystem through which requests and confirmations are routed. See SAM-3.

3.1.39 SCSI target port: A SCSI target device object that contains a task router and acts as the connection between device servers and task managers and the service delivery subsystem through which indications and responses are routed. See SAM-3.

3.1.40 secret key: A value that is known to only a limited set of at least two entities (e.g., the device server and security manager) and serves as input for an integrity check value (see 3.1.19) computation.

3.1.41 security manager: The component of an OSD configuration (see 4.4) that manages secret keys (see 3.1.40) and prepares secure credentials (see 3.1.11) containing capabilities (see 3.1.4) thus granting application clients specified access to a specified OSD logical unit. See 4.12.

3.1.42 security method: A set of zero or more security features and algorithms from the OSD security model that are enabled as a group to thwart zero or more security threats. See 4.12.4.

3.1.43 security token: A value representing an I_T nexus (see 3.1.18) known to both the application client and device server. See 4.12.4.2.

3.1.44 sense data: Data describing an error or exceptional condition that a device server delivers to an application client. See SPC-3.

3.1.45 sense key: The contents of the SENSE KEY field in the sense data (see 3.1.44).

3.1.46 solo capability: A capability (see 3.1.4) that is the first capability returned in a credential (see 3.1.11) and is placed in the CDB (see 5.2.1). A credential may bind a solo capability to one or more extension capabilities (see 3.1.16).

3.1.47 space-padded: A type of field in which unused bytes are filled with ASCII space (20h) characters. See 3.8.2.

3.1.48 stable storage: Storage that survives all the events that may result in the loss of data in the volatile cache (see 3.1.57). See 4.14.

3.1.49 status: One byte of response information sent from a device server to an application client upon completion of each command. See SAM-3.

3.1.50 task: A SCSI architecture model object within a logical unit that represents the work associated with a command. See SAM-3.

3.1.51 undefined attribute: An attribute (see 3.1.3) for which the only attribute length known to the OSD logical unit (see 3.1.28) is zero (i.e., an attribute whose attribute page and attribute number are valid but that is unknown to the OSD logical unit).

3.1.52 universal time (UT): The time at longitude zero, colloquially known as Greenwich Mean Time. See <http://aa.usno.navy.mil/faq/docs/UT.html>.

3.1.53 user object: An OSD object (see 3.1.29) that contains user data (see 4.6.1) that is referenced by byte offset within the OSD object.

3.1.54 User_Object_ID: The identifier for one user object (see 3.1.53).

3.1.55 user tracking collection: A TRACKING collection (see 4.6.6.3) with a Collection_Object_ID that is greater than or equal to 10000h (i.e., a TRACKING collection that is not a well-known collection (see 4.6.6.5)).

3.1.56 vendor specific (VS): Something (e.g., a bit, field, code value, behavior) that is not defined by this standard and may be vendor defined.

3.1.57 volatile cache: Storage that is lost after a power on or hard reset event (see SAM-4) and may be lost after an I_T nexus loss or logical unit reset event (see SAM-4). See 4.14.

3.1.58 zero-padded: A type of field in which unused bytes are filled with zeros. See 3.8.2.

3.2 Acronyms

+	added to
÷	divided by
×	multiplied by
C	A constant equal to 6000 0000h used in describing object attribute page numbers (see 4.8.5)
CDB	Command Descriptor Block (see 3.1.9)
DH	Diffie-Hellman (see 2.4, RFC 2409 and RFC 3526)
FIPS	Federal Information Processing Standard (see 2.3)
HMAC-SHA1	Keyed-Hash Message Authentication Code - Secure Hash Algorithm 1 (see 2.3)
I/O	Input/Output
IANA	Internet Assigned Numbers Authority (see http://www.iana.org)
ID	Identifier
INCITS	InterNational Committee for Information Technology Standards
ISO	Organization for International Standards
LSB	Least Significant Bit
MODP	Modular Exponential (see 2.4, RFC 3526)
MSB	Most Significant Bit
n/a	not applicable
OBSD	An Object-Based Storage Device, a SCSI device that implements this standard (see 3.1.27)
OSD	Object-based Storage Device Commands (see 2.2)
OSD-2	Object-based Storage Device Commands -2 (this standard, see clause 1)
P	A constant equal to 3000 0000h used in describing object attribute page numbers (see 4.8.5)
R	A constant equal to 9000 0000h used in describing object attribute page numbers (see 4.8.5)
RAID	Redundant Array of Independent Disks
SAM-3	SCSI Architecture Model -3 (see 2.1)
SAM-4	SCSI Architecture Model -4 (see 2.5)
SAN	Storage Area Network (see Storage Networking Industry Association web site, www.snia.org)
SBC	SCSI-3 Block Commands (see clause 1)

SCSI	The architecture defined by the family of standards described in clause 1
SPC-3	SCSI Primary Commands -3 (see 2.1)
SPC-4	SCSI Primary Commands -4 (see 2.5)
UT	Universal Time (see 3.1.52)
VPD	Vital Product Data (see SPC-3)
VS	Vendor Specific (see 3.1.56)

3.3 Keywords

3.3.1 expected: A keyword used to describe the behavior of the hardware or software in the design models assumed by this standard. Other hardware and software design models may also be implemented.

3.3.2 ignored: A keyword used to describe an unused bit, byte, word, field or code value. The contents or value of an ignored bit, byte, word, field or code value shall not be examined by the receiving SCSI device and may be set to any value by the transmitting SCSI device.

3.3.3 invalid: A keyword used to describe an illegal or unsupported bit, byte, word, field or code value. Receipt of an invalid bit, byte, word, field or code value shall be reported as an error.

3.3.4 mandatory: A keyword indicating an item that is required to be implemented as defined in this standard.

3.3.5 may: A keyword that indicates flexibility of choice with no implied preference (equivalent to "may or may not").

3.3.6 may not: A keyword that indicates flexibility of choice with no implied preference (equivalent to "may or may not").

3.3.7 obsolete: A keyword indicating that an item was defined in prior SCSI standards but has been removed from this standard.

3.3.8 optional: A keyword that describes features that are not required to be implemented by this standard. However, if any optional feature defined by this standard is implemented, then it shall be implemented as defined in this standard.

3.3.9 prohibited: A keyword used to describe a feature, function, or coded value that is defined in a non-SCSI standard (i.e., a standard that is not a member of the SCSI family of standards) to which this standard makes a normative reference where the use of said feature, function, or coded value is not allowed for implementations of this standard.

3.3.10 reserved: A keyword referring to bits, bytes, words, fields and code values that are set aside for future standardization. A reserved bit, byte, word or field shall be set to zero, or in accordance with a future extension to this standard. Recipients are not required to check reserved bits, bytes, words or fields for zero values. Receipt of reserved code values in defined fields shall be reported as an error.

3.3.11 restricted: A keyword referring to bits, bytes, words, and fields that are set aside for use in other SCSI standards. A restricted bit, byte, word, or field shall be treated as a reserved bit, byte, word or field for the purposes of the requirements defined in this standard.

3.3.12 shall: A keyword indicating a mandatory requirement. Designers are required to implement all such mandatory requirements to ensure interoperability with other products that conform to this standard.

3.3.13 should: A keyword indicating flexibility of choice with a strongly preferred alternative; equivalent to the phrase "it is strongly recommended".

3.3.14 x or xx: The value of the bit or field is not relevant.

3.4 Conventions

Certain words and terms used in this standard have a specific meaning beyond the normal English meaning. These words and terms are defined either in 3.1 or in the text where they first appear. Names of commands, statuses, sense keys, and additional sense codes are in all uppercase (e.g., IDENTIFY DEVICE). Lowercase is used for words having the normal English meaning.

The names of fields are in small uppercase (e.g., STARTING BYTE ADDRESS). Normal case is used when the contents of a field are being discussed. Fields containing only one bit are usually referred to as the NAME bit instead of the NAME field.

A binary number is represented in this standard by any sequence of digits consisting of only the Western-Arabic numerals 0 and 1 immediately followed by a lower-case b (e.g., 0101b). Underscores or spaces may be included in binary number representations to increase readability or delineate field boundaries (e.g., 0 0101 1010b or 0_0101_1010b).

A hexadecimal number is represented in this standard by any sequence of digits consisting of only the Western-Arabic numerals 0 through 9 and/or the upper-case English letters A through F immediately followed by a lower-case h (e.g., FA23h). Underscores or spaces may be included in hexadecimal number representations to increase readability or delineate field boundaries (e.g., B FD8C FA23h or B_FD8C_FA23h).

A decimal number is represented in this standard by any sequence of digits consisting of only the Western-Arabic numerals 0 through 9 not immediately followed by a lower-case b or lower-case h (e.g., 25).

When the value of the bit or field is not relevant, x or xx appears in place of a specific value.

A range of numeric values is represented in this standard in the form "a to z", where a is the first value included in the range, all values between a and z are included in the range, and z is the last value included in the range (e.g., the representation "0h to 3h" includes the values 0h, 1h, 2h, and 3h).

This standard uses the following conventions for representing decimal numbers:

- a) The decimal separator (i.e., separating the integer and fractional portions of the number) is a period;
- b) The thousands separator (i.e., separating groups of three digits in a portion of the number) is a space; and
- c) The thousands separator is used in both the integer portion and the fraction portion of a number.

Table 1 shows some examples of decimal numbers represented using various conventions.

Table 1 — Numbering conventions examples

French	English	This standard
0,6	0.6	0.6
3,141 592 65	3.14159265	3.141 592 65
1 000	1,000	1 000
1 323 462,95	1,323,462.95	1 323 462.95

A decimal number represented in this standard with an overline over one or more digits following the decimal point is a number where the overlined digits are infinitely repeating (e.g., $666.\overline{6}$ means $666.666\ 666\dots$ or $666\ 2/3$ and $12.\overline{142\ 857}$ means $12.142\ 857\ 142\ 857\dots$ or $12\ 1/7$).

Lists sequenced by letters (e.g., a-red, b-blue, c-green) show no priority relationship between the listed items. Numbered lists (e.g., 1-red, 2-blue, 3-green) show a priority ordering between the listed items.

If a conflict arises between text, tables, or figures, the order of precedence to resolve the conflicts is text; then tables; and finally figures. Not all tables or figures are fully described in the text. Tables show data format and values. Notes do not constitute any requirements for implementors.

3.5 Bit and byte ordering

This subclause describes the representation of fields in a table that defines the format of a SCSI structure (e.g., the format of a CDB).

If a field consists of more than one bit and contains a single value (e.g., a number), the least significant bit (LSB) is shown on the right and the most significant bit (MSB) is shown on the left (e.g., in a byte, bit 7 is the MSB and is shown on the left; and bit 0 is the LSB and is shown on the right). The MSB and LSB are not labeled if the field consists of 8 or fewer bits.

If a field consists of more than one byte and contains a single value, the byte containing the MSB is stored at the lowest address and the byte containing the LSB is stored at the highest address (i.e., big-endian byte ordering). The MSB and LSB are labeled.

If a field consists of more than one byte and contains multiple fields each with their own values (e.g., a descriptor), there is no MSB and LSB of the field itself and thus there are no MSB and LSB labels. Each individual field has an MSB and LSB that are labeled as appropriate in the table (if any) that describes the format of the sub-structure having multiple fields.

If a field contains a text string (e.g., ASCII), the MSB label is the MSB of the first character and the LSB label is the LSB of the last character.

When required for clarity, multiple byte fields may be represented with only two rows in a table. This condition is represented by values in the byte number column not increasing by one in each subsequent table row, thus indicating the presence of additional bytes.

3.6 Signed and unsigned integers values

Unless otherwise stated, all values defined by this standard are transferred as unsigned integers. Signed integers are:

- a) Positive valued unsigned integers when the most significant bit is set to zero; and
- b) Negative valued integers in two's complement form (e.g., in a one-byte field, a negative valued one is represented as FFh) when the most significant bit is set to one.

3.7 Notation conventions

3.7.1 Notation for byte encoded character strings

When this standard requires one or more bytes to contain specific encoded characters, the specific characters are enclosed in double quotation marks. The double quotation marks identify the start and end of the characters that are required to be encoded but the quotation marks are not to be encoded. The characters that are to be encoded are shown in exactly the case that is to be encoded.

The encoded characters and the double quotation marks that enclose them are preceded by text that specifies the character encoding methodology and the number of characters required to be encoded.

Using the notation described in this subclause, stating that eleven ASCII characters "SCSI device" are to be encoded would be the same as writing out the following sequence of byte values: 53h 43h 53h 49h 20h 64h 65h 76h 69h 63h 65h.

3.7.2 Notation for procedure calls

In this standard, the model for functional interfaces between objects is a procedure call. Such interfaces are specified using the following notation:

[Result =] Procedure Name (IN ([input-1] [,input-2] ...)), OUT ([output-1] [,output-2] ...))

Where:

Result: A single value representing the outcome of the procedure call.

Procedure Name: A descriptive name for the function modeled by the procedure call. When the procedure call model is used to describe a SCSI transport protocol service, the procedure name is the same as the service name.

Input-1, Input-2, ...: A comma-separated list of names identifying caller-supplied input arguments.

Output-1, Output-2, ...: A comma-separated list of names identifying output arguments to be returned by the procedure call.

[...]: Brackets enclosing optional or conditional arguments.

This notation allows arguments to be specified as inputs and outputs. The following is an example of a procedure call specification:

Found = Search (IN (Pattern, Item List), OUT ([Item Found]))

Where:

Found = Flag

Flag: if set to one, indicates that a matching item was located.

Input Arguments:

Pattern = ... /* Definition of Pattern argument */
Argument containing the search pattern.

Item List = Item<NN> /* Definition of Item List as an array of NN Item arguments*/
Contains the items to be searched for a match.

Output Arguments:

Item Found = Item ... /* Item located by the search procedure call */
This argument is only returned if the search succeeds.

3.8 Data field requirements

3.8.1 ASCII data field requirements

ASCII data fields shall contain only ASCII graphic codes (i.e., code values 20h through 7Eh) and may be terminated with one or more ASCII null (00h) characters.

3.8.2 Data field termination and padding requirements

A data field that is described as being null-terminated shall have one byte containing an ASCII null (00h) character in the last used byte (i.e., highest offset) of the field and no other bytes in the field shall contain the ASCII null character.

A data field may be specified to be a fixed length that may be larger than the contents need or a data field may be specified to have a length that is a multiple of a given value (e.g., a multiple of four bytes).

When such fields are described as being space-padded, the bytes at the end of the field that are not needed to contain the field data shall contain ASCII space (20h) characters.

When such fields are described as being null-padded, the bytes at the end of the field that are not needed to contain the field data shall contain ASCII null (00h) characters.

When such fields are described as being zero-padded, the bytes at the end of the field that are not needed to contain the field data shall contain zeros.

NOTE 1 - There is no difference between the pad byte contents in null-padded and zero-padded fields. The difference is in the format of the other bytes in the field.

A data field that is described as being both null-terminated and null-padded shall have at least one byte containing an ASCII null (00h) character in the end of the field (i.e., highest offset) and may have more than one byte containing ASCII null characters if needed to meet the specified field length requirements. If more than one byte in a null-terminated, null-padded field contains the ASCII null character, all the bytes containing the ASCII null character shall be at the end of the field (i.e., only the highest offsets).

4 SCSI OSD Model

4.1 The request-response model

The SCSI command set assumes an underlying request-response protocol. The fundamental properties of the request-response protocol are defined in SAM-4. Action on OSD commands shall not be deemed completed until a response is received. The response shall include a status that indicates the final disposition of the command. As per SAM-4, the request-response protocol may be modeled as a procedure call, specifically:

Service response = Execute Command (IN (I_T_L_Q Nexus, CDB, Task Attribute, [Data-In Buffer Size],
[Data-Out Buffer], [Data-Out Buffer Size], [Command Reference Number], [Task Priority]),
OUT ([Data-In Buffer], [Sense Data], [Sense Data Length], Status))

SAM-4 defines all of the inputs and outputs in the procedure call above. As they apply to an OBSD (see 3.1.27), this standard defines the contents of the following procedure inputs and outputs; CDB, Data-Out Buffer, Data-Out Buffer Size, Data-In Buffer, Data-In Buffer Size, and Sense Data. This standard does not define all possible instances of these procedure inputs and outputs. This standard defines only those instances that apply to an OBSD.

This standard references values returned via the Status output parameter (e.g., GOOD and CHECK CONDITION). Status values are not defined by this standard. SAM-4 defines all Status values.

The entity that makes the procedure call is an application client (see SAM-4). The procedure call's representation arrives at the SCSI target device in the form of a device service request. The entity that performs the work of the procedure call is a device server (see SAM-4).

4.2 OSD type devices

An OBSD (see 3.1.27) contains one or more logical units that return the OSD peripheral device type value in response to an INQUIRY command (see SPC-3). From the perspective of the application client, an OBSD logical unit contains OSD objects (see 3.1.29), not logical blocks (see 4.5). All stored data objects (see 4.6) have associated attributes (see 4.8).

4.3 OSD object abstraction

The OSD object abstraction is designed to re-divide the responsibility for managing the access to data on a storage device by assigning to the storage device additional responsibilities in the area of space management. Figure 2 shows the relationship between the OSD model and a traditional SBC-based model for a file system.

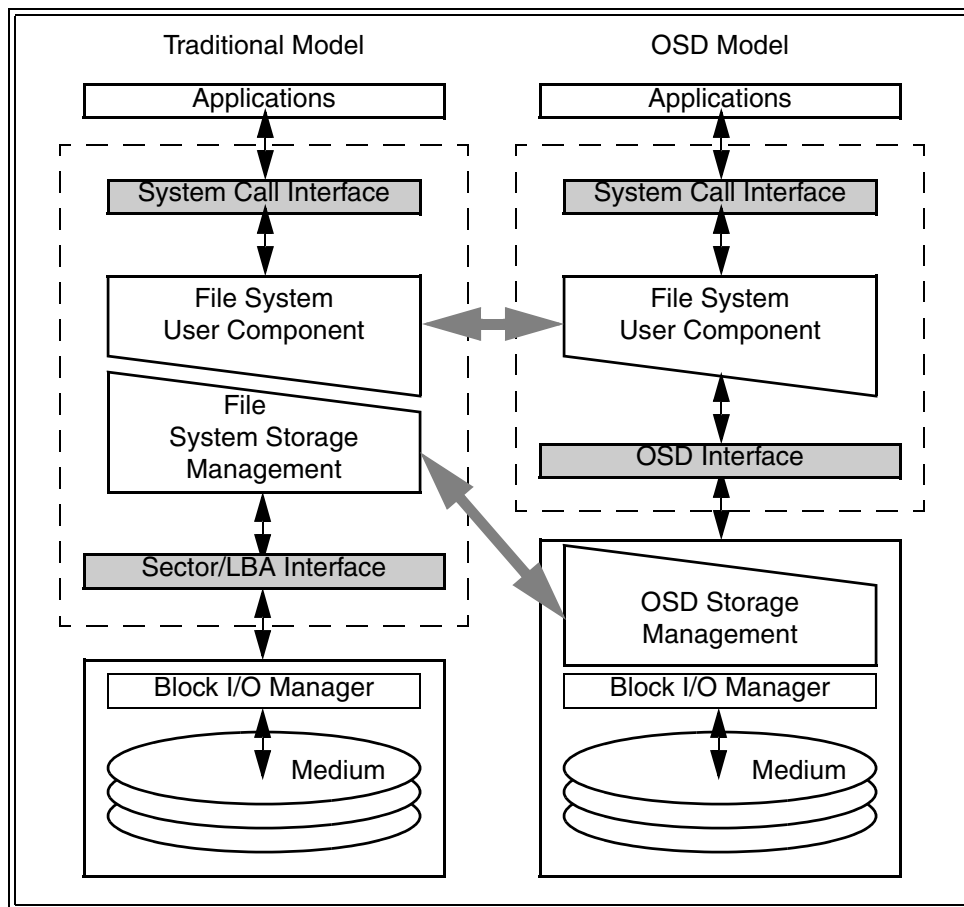


Figure 2 — Comparison of traditional and OSD storage models

The user component of the file system contains such functions as:

- a) Hierarchy management;
- b) Naming; and
- c) User access control.

The storage management component is focused on mapping logical constructs (e.g., files or database entries) to the physical organization of the storage media. In the OSD model, the logical constructs are called user objects (see 4.6.5). The root object (see 4.6.3), partitions (see 4.6.4), and collections (see 4.6.6) provide additional navigational aids for user objects.

In addition to mapping data, the storage management component maintains other information about the OSD objects that it stores (e.g., size, and usage quotas, and associated username) in attributes (see 4.8). The user component may have the ability to influence the properties of object data through the specification of attributes (e.g., directing that the location of an object to be in close proximity to another object or to have some higher performance characteristic) via mechanisms that are outside the scope of this standard.

In this model, the OBSD (see 3.1.27) makes the decisions as to where to allocate storage capacity for individual data entities and managing free space.

4.4 Elements of the example configuration

The example in this subclause (see figure 3) illustrates the three mandatory and two optional constituents of an OSD configuration:

- a) Object-Based Storage Devices;
- b) Service delivery subsystem;
- c) Host systems (i.e., initiator devices);
- d) Optionally, a security manager; and
- e) Optionally, a policy/storage manager.




Figure 3 — Example OSD Configuration

The OBSDs are the storage components of the system to be shared (i.e., disc drives, RAID subsystems, tape drives, tape libraries, optical drives, jukeboxes, or other storage devices).

Application clients using multiple SCSI initiator ports share directly access an OBSD (see 3.1.27) via the service delivery subsystem. The service delivery subsystem is used by the components in the OSD model, except possibly policy/storage manager and/or the security manager, to intercommunicate. The OSD security model (see 4.12) does not require the service delivery subsystem to provide security-related services (i.e., authentication and confidentiality), but is designed to take advantage of whatever security-related services are provided.

The policy/storage manager (see 4.11), if present, coordinates access constraints between OSD device servers and application clients, preparing the capabilities (see 3.1.4) application clients place in CDBs or CDB continuation segments (see 5.3) to gain access to OSD objects and command functions.

The security manager (see 4.12), if present, secures capabilities in cryptographically protected credentials (see 3.1.11) for OSD device servers and application clients.

The policy/storage manager and security manager may reside in the OBSDs, in application clients, or as separate entities.

The policy/storage manager and security manager may use the service delivery subsystem and be an application client, but they also may use another mechanism to communicate with the OSD device servers and application clients. Security-related requirements on the communications mechanisms used by the security manager are described in 4.12.2.

4.5 Description of the OSD Architecture

Data is stored in abstract containers by the OBSD (see 3.1.27) logical unit. Data in the abstract containers is not addressable using LBAs (Logical Block Addresses). The OSD logical unit allocates space for data and delivers a unique identifier to the application client. The application client uses the same unique identifier for subsequent accesses to the data.

In addition to the objects defined in SAM-3, this standard provides the OSD model objects listed in table 2.

Table 2 — OSD model objects

OSD model objects representing stored data		OSD model objects representing transient application client activities	
OSD Object	Reference	OSD Object	Reference
Root Object	4.6.3	Capability	4.11.2
Partition	4.6.4	Credential	4.12.5.1
Collection	4.6.6		
User Object	4.6.5		
Associated Data	Reference		
Attributes	4.8		

4.6 Stored data objects

4.6.1 Stored data object types

An OBSD contains one or more logical units with the following types of stored data objects:

- a) **Root object:** Each OSD logical unit contains one and only one root object. Its attributes (see 4.8) contain global characteristics for the OSD logical unit (e.g., the total capacity of the logical unit and number of partitions that it contains). The root object contains a list of Partition_IDs for the partitions in the logical unit that may be retrieved using the LIST command (see 6.20).
- b) **Partition:** This OSD object is created by specific commands from an application client. It contains a set of collections and user objects that share common security requirements and attributes (e.g., the default security method and a capacity quota). The default values for some partition attributes are copied from specified attributes in the root object. Each partition contains a list of User_Object_IDs and Collection_Object_IDs contained in the partition that may be retrieved using the LIST command (see 6.20) and LIST COLLECTION command (see 6.21) command, respectively.
- c) **Collection:** This OSD object is created by commands from an application client. It is used for fast indexing of user objects. A collection is contained within one partition. A partition may contain zero or more collections. A user object may be a member of zero or more collections concurrently. Support for collections is optional. Default values for some collection attributes are copied from specified attributes of the partition in which it is listed. Each collection contains a list of User_Object_IDs contained in the collection that may be retrieved using the LIST COLLECTION command (see 6.21).
- d) **User object:** This OSD object contains end-user data (e.g., file or database data). Its attributes include the logical size of the user data and timestamps for creation, access, and modification of the end user data. Default values for some user object attributes are copied from specified attributes of the partition in which it is listed.

An OSD logical unit shall always contain a root object and an OSD object for partition zero (see 3.1.33) with at least the attributes (see 4.8) defined by this standard.

4.6.2 Identifying OSD objects

The combination of Partition_ID and User_Object_ID uniquely identifies the root object, each partition, each collection, and each user object. Partition_ID and User_Object_ID values are assigned as shown in table 3.

Table 3 — Partition_ID and User_Object_ID value assignments

Partition_ID	User_Object_ID or Collection_Object_ID	Description
0h	0h	Root object
0h	1h to FFFF FFFF FFFF FFFFh	Reserved
1h to FFFFh	0h to FFFF FFFF FFFF FFFFh	Reserved
10000h to FFFF FFFF FFFF FFFFh	0h	Partition ^a
	1h to 0FFFh	Reserved
	1000h to BFFFh	Well known collections ^b
	C000h to FFFFh	Reserved
10000h to FFFF FFFF FFFF FFFFh	10000h to FFFF FFFF FFFF FFFFh	Collection or User object ^c
^a Partition_ID values assigned by the OSD logical unit in response to application client requests. ^b Well known collections have constant Collection_Object_ID values and may be members of any partition (see 4.6.6.5). ^c Collection_Object_ID values and User_Object_ID values assigned by the OSD logical unit in response to application client requests.		

4.6.3 Root object

There is one root object per OSD logical unit. The root object is addressed by setting both Partition_ID value and User_Object_ID value to zero. The root object is the starting point for navigation of the structure on an OSD logical unit.

The root object does not contain a read/write data area. The device server shall terminate all READ commands, WRITE commands, and APPEND commands sent to the root object with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

4.6.4 Partitions

User objects are collected into partitions, that are represented by partition OSD objects. There may be any number of partitions, up to a specified quota or the capacity of the OSD logical unit.

A Partition_ID uniquely identifies each partition. Partitions have a User_Object_ID of zero and a Partition_ID (see 4.6.2) that is assigned by the OSD logical unit when the partition is created. The partition with Partition_ID zero represents the root object and is called partition zero.

When a partition is created using the CREATE PARTITION command (see 6.9), a partition OSD object shall be created to provide navigation among user objects in the partition.

To obtain a list of the valid Partition_IDs, an application client sends the LIST command (see 6.20) to the device server specifying the root object.

A partition does not contain a read/write data area. The device server shall terminate all READ commands, WRITE commands, and APPEND commands sent to a partition with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

4.6.5 User objects

User objects contain end-user data (i.e., the content of this data is owned by the applications that cause the creation, writing and reading the user objects). User objects have the Partition_ID of the partition to which they belong and a User_Object_ID (see 4.6.2) that is assigned by the OSD logical unit when the user object is created. A user object is a member of only one partition.

Within a single partition, no user object shall be assigned the same User_Object_ID value as any Collection_Object_ID and no collection shall be assigned the same Collection_Object_ID as any User_Object_ID (i.e., collections and user objects share the same number space for their identifier values).

A user object may be made a member of one or more collections (see 4.6.6) by setting attribute values in the user object's Collections attributes page (see 7.1.3.21).

4.6.6 Collections

4.6.6.1 Overview

Support for collections is optional. If collections are not supported:

- a) The length of attribute number 4h in the User Object Directory attributes page (see 7.1.3.7) shall be zero for every user object (i.e., no Collections attributes pages identified); and
- b) Zero shall be returned as the length of attribute number 0h in every Collections attributes page (see 7.1.3.21).

A partition may contain zero or more collections each of which may contain zero or more user objects. One user object may be a member of zero or more collections.

Collections have the Partition_ID of the partition to which they belong and a Collection_Object_ID (see 4.6.2) that is assigned by the OSD logical unit when the collection is created. A collection is a member of only one partition.

Within a single partition, no collection shall be assigned the same Collection_Object_ID as any User_Object_ID and no user object shall be assigned the same User_Object_ID value as any Collection_Object_ID (i.e., collections and user objects share the same number space for their identifier values).

The LIST COLLECTION command (see 6.21) lists all the collections in a partition or all the user objects that are members of a collection.

A collection does not contain a read/write data area. The device server shall terminate all READ commands, WRITE commands, and APPEND commands sent to the collection with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The collection type attribute in the Collection Information attributes page indicates the type of the collection as shown in table 171 (see 7.1.3.10). Different collection types have different operational characteristics. The following collection types are defined:

- a) LINKED (see 4.6.6.2);
- b) TRACKING (see 4.6.6.3); and
- c) SPONTANEOUS (see 4.6.6.4).

4.6.6.2 LINKED collections

The device server maintains a linkage between the user object entries in a LINKED collection and the actual user objects (e.g., if a user object that is a member of a LINKED collection is removed from the partition, all entries for the user object are removed from all LINKED collections of which the user object is a member).

A LINKED collection is created using the CREATE COLLECTION command (see 6.8) and deleted using the REMOVE COLLECTION command (see 6.32). The page format of the Collections attributes page (see 7.1.3.21) lists all the collections in which a user object is a member.

User objects are added to or removed from the membership of a LINKED collection by setting attribute values in the user object's Collections attributes page (see 7.1.3.21).

4.6.6.3 TRACKING collections

The members of a TRACKING collection have no linkage to the actual user objects (e.g., removal of a user object from the partition or replacement of a user object with another user object having the same User_Object_ID have no effect on the user object's membership in the TRACKING collection).

TRACKING collections are used to track the progress of commands that operate on multiple objects (e.g., multi-object commands (see 4.6.6.6) and the CREATE SNAPSHOT command (see 6.10)).

The Collection_Object_ID of a TRACKING collection affects the collection's operational characteristics as shown in table 4.

Table 4 — TRACKING collection operational characteristics

Collection_Object_ID	Allowed members	Created by	Removed by	Dynamic addition of members allowed
0h	Reserved (see table 3 in 4.6.2)			
1h to FFFFh	user objects and collections	The device server in response to specific commands (e.g., the CREATE SNAPSHOT command (see 6.10))		Yes
10000h to FFFF FFFF FFFF FFFFh	user objects	A CREATE USER TRACKING COLLECTION command (see 6.11)	A REMOVE COLLECTION command (see 6.32)	No

TRACKING collections shall include the Command Tracking attributes page (see 7.1.3.20).

4.6.6.4 SPONTANEOUS collections

All SPONTANEOUS collections are well known collections (see 4.6.6.5). The SPONTANEOUS collection's Collection_Object_ID specifies how to determine the collection's membership. A SPONTANEOUS collection's membership is recomputed each time it is retrieved.

4.6.6.5 Well known collections

4.6.6.5.1 Overview

Any partition except partition zero (see 3.1.33) may contain one or more of the well known collections shown in table 5.

Table 5 — Well known collections

Collection_ Object_ID ^a	Description	Type	Support Requirements	Reference
1000h to 1081h	Reserved			
1082h	All user objects in partition	SPONTANEOUS ^b	Mandatory	4.6.6.5.2
1083h to 8000h	Reserved			
8001h	Snapshot/clone tracking	TRACKING ^c	Optional	4.6.6.5.3
8002h to BFFFh	Reserved			
^a Collection_Object_IDs not shown in this table are shown in table 3 (see 4.6.2). ^b See 4.6.6.4. ^c See 4.6.6.3.				

Well known collections are not included in the parameter data returned by a LIST COLLECTION command (see 6.21).

4.6.6.5.2 The all user objects in partition well known collection

The membership of the all user objects in partition well known collection is all user objects in the partition which contains it. The effect of each access to the all user objects in partition well known collection is the equivalent of:

- 1) Processing a LIST command (see 6.20) with:
 - A) The PARTITION_ID field set to the Partition_ID of the partition that contains the all user objects in partition well known collection; and
 - B) The LIST_ATTR bit set to zero;
 and
- 2) Defining the collection's membership to match the command's output.

4.6.6.5.3 The snapshot/clone tracking well known collection

The snapshot/clone tracking well known collection is a member of any partition that is:

- a) The destination partition for a:
 - A) CREATE SNAPSHOT command (see 6.10);
 - B) CREATE CLONE command (see 6.7); or
 - C) REFERESH SNAPSHOT OR CLONE command (see 6.30);
 or
- b) The source partition for a RESTORE PARTITION FROM SNAPSHOT command (see 6.35).

The membership and attributes of a snapshot/clone tracking well known collection provide sufficient information to:

- a) Track the progress of the creation of the snapshot or clone partition; and
- b) Restart an interrupted CREATE SNAPSHOT command, CREATE CLONE command, REFERESH SNAPSHOT OR CLONE command, or RESTORE PARTITION FROM SNAPSHOT command.

The device server shall update the snapshot/clone tracking well known collection and continue processing the operation that it describes until the operation is completed with or without an error regardless of the conditions detected during such processing (e.g., reset events or power on events described in SAM-4).

4.6.6.6 Commands that use collections to affect multiple user objects

Commands such as SET MEMBER ATTRIBUTES (see 6.39) (i.e., multi-object commands) process multiple user objects using the membership of a collection that is not a well known collection (see 4.6.6.5) as a list of the user objects on which the specified operations are to be performed.

With the exception of the REMOVE MEMBER OBJECTS command (see 6.33), multi-object commands process only user tracking collections (see 3.1.55). If the COLLECTION_OBJECT_ID field in a multi-object command CDB other than a REMOVE MEMBER OBJECTS command specifies a collection for which the collection type attribute in the Collection Information attributes page (see 7.1.3.10) contains a value other TRACKING, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the COLLECTION_OBJECT_ID field in a multi-object command CDB specifies a object that is not a collection or is a well known collection (see 4.6.6.5), then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the COLLECTION_OBJECT_ID field in a multi-object command CDB specifies a user tracking collection in which the active command status attribute in the Command Tracking attributes page (see 7.1.3.20) is not set to zero, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

Before any user object is processed, the attributes in the Command Tracking attributes page (see 7.1.3.20) shall be updated to reflect active use of the collection for processing the command.

Each user object in the specified collection shall be processed as follows:

- a) If the user object has been removed, the specified operation shall not be performed on the non-existent user object. This shall not be considered to be an error;
- 2) If the user object has not been removed and the creation time attribute in the User Object Timestamps attributes page (see 7.1.3.18) is earlier than or equal to (i.e., less than or equal to) the creation time attribute in the Collection Timestamps attributes page (see 7.1.3.17) (i.e., if the user object has not been replaced), then the quotas (see 4.10) that apply to the specified operation shall be evaluated and processing of the operation shall be handled as follows:
 - A) If a quota error condition is detected, the multi-object command shall be terminated as described in this subclause; or
 - B) If no quota error condition is detected, the specified operation shall be performed on the user object and whether or not an error is detected shall be noted;
- 3) If the creation time attribute in the User Object Timestamps attributes page is later than (i.e., greater than) the creation time attribute in the Collection Timestamps attributes page, then the specified operation shall not be performed on the user object. This shall not be considered to be an error;
- 4) If no error has been detected, the user object shall be removed from the specified collection; and
- 5) The attributes in the Command Tracking attributes page shall be updated to reflect completion of processing for the user object.

As a result of these requirements, the following conditions apply:

- a) After an error condition that prevented processing of all user objects in the collection is corrected, the same command specifying the same collection may be sent to continue processing; and
- b) Application clients may poll to determine the progress of a multi-object command using the LIST COLLECTION command (see 6.21) and/or the contents of the Command Tracking attributes page (see 7.1.3.20).

NOTE 2 The LIST command and LIST COLLECTION command are not multi-object commands.

Two multi-object commands shall not concurrently process the same collection. If a multi-object command is received with the COLLECTION_OBJECT_ID field in the CDB specifying the Collection_Object_ID (see 4.6.2) of a collection that is already being processed by a different multi-object command, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The device server may process more than one user object concurrently.

If an error is detected during the processing of a user object:

- a) The user object shall not be removed from the collection;
- b) Processing that has already been started on any other user object shall be completed to the greatest degree possible and any user objects for which processing is successfully completed shall be removed from the collection;
- c) If necessary, the policy access tag attribute in the User Object Policy/Security attributes page (see 7.1.3.25) for any user object for which an error is detected shall be updated as described in 4.11.3.2;
- d) Processing shall not be started for any user object that has not already started processing; and
- e) When no user objects are being processed:
 - A) If the IMMED_TR bit (see 5.2.5) is set to zero, the command shall be terminated with the status and sense data corresponding to the first error that was detected; and
 - B) The attributes in the Command Tracking attributes page (see 7.1.3.20) shall be updated to reflect the termination of processing, including the status and additional sense data with which the command, if any would have been or was terminated.

If a multi-object command is terminated as part of processing a command-related condition (see SAM-4), a task management function, or as the result of a SCSI device condition (e.g., logical unit reset) established in response to an event (see SAM-4), then the device server shall:

- a) Update the attributes in the Command Tracking attributes page (see 7.1.3.20) reflect the interruption of the command; and
- b) Either:
 - A) Establish a consistent, stable state for each user object being processed; or
 - B) Set the policy access tag attribute in the User Object Policy/Security attributes page described in 4.11.3.2 for any user object for which it is not possible to establish consistent state.

The device server shall not remove the specified collection upon completion of the multi-object command, even if the collection contains zero user objects.

If the CDB GET/SET CDBFMT field contains 11b (i.e., when list format attributes processing is specified), multi-object commands allow setting and retrieving of both collection attributes and user object attributes. The get and set

attributes parameters are described in 5.2.6. The list format is described in 7.1.4. Attribute retrieval and setting shall be processed as shown in table 6.

Table 6 — Attributes retrieval and setting requirements for multi-object commands

Attribute page number values	Command	Description
C+0h to C+2FFF FFFFh and F000 0000h to FFFF FFFFh	Any multi-object command	The attribute values shall be returned in the retrieved attributes segment of the Data-In Buffer (see 4.15.3) as described in 5.2.6.4 using list type Eh (see 7.1.4.4). The setting of attributes shall be processed as described in 5.2.6.4.
0h to 2FFF FFFFh	GET MEMBER ATTRIBUTES or SET MEMBER ATTRIBUTES	The attribute values for every user object that is a member of the collection shall be returned in the retrieved attributes segment of the Data-In Buffer as described in 5.2.6.4 using list type Eh (see 7.1.4.4). The setting of attributes shall be processed as described in 5.2.2.3 and the same user object attribute values shall be set in every user object that is a member of the collection.
	All other multi-object commands	The command to be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.
All other page number values	Any multi-object command	The command to be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

Multi-object commands may allow retrieval and setting of attribute values using page format (see 7.1.3), but only collection attribute pages shall be processed in page format.

The multi-object operation in progress attribute in the Collection Information attributes page (see 7.1.3.10) shall be set as follows:

- a) To one before an operation is performed as described in this subclause on the first user object in a collection; and
- b) To zero before the processing of a multi-object command is completed or terminated as described in this subclause.

4.7 Data object accessibility

Write access to user object data (see 4.6.5), collection membership (see 4.6.6), the creation of new user object, the creation of new collections, the creation of new partitions, or the attributes (see 4.8) associated with any type of object, may be controlled using:

- a) Policy/storage manager capabilities (see 4.11.2); or
- b) Object accessibility attributes in the following attributes pages:
 - A) The User Object Information attributes page (see 7.1.3.11);
 - B) The Collection Information attributes page (see 7.1.3.10);
 - C) The Partition Information attributes page (see 7.1.3.9); and

D) The Root Information attributes page (see 7.1.3.8).

The object accessibility attributes form the following prioritized hierarchy:

- 1) Root (i.e., the root object, all partitions, all collections, and all user objects);
- 2) Partition (i.e., all partition attributes, all collections, and all user objects); and
- 3) A leaf object (i.e., one collection or one user object).

Denial of write access at a higher level in the hierarchy includes denial of access in all lower levels (e.g., write access to a collection depends on write access being allowed for the collection, the partition, and the root object).

Denial of write access to an object with members means denial of the ability to create new members in that object (e.g., denial of write access to the root object means denial of the ability to create new partitions).

If a command attempts a write access that is not allowed by all applicable levels of the object accessibility attributes hierarchy, then:

- a) No part of the command shall be processed; and
- b) The command shall be terminated with CHECK CONDITION status, with the sense key set to DATA PROTECT, the additional sense code set to CONDITIONAL WRITE PROTECT, and the INFORMATION field set as shown in table 7.

Table 7 — DATA PROTECT INFORMATION field contents

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
5								
6	ATTRIBUTE	Reserved						
7	OBJECT TYPE							

If the requested write access that is denied is to the data in a user object, the membership list in a collection, the membership list in a partition, or the membership list in the root object, then the ATTRIBUTE bit shall be set to zero. If the requested write access that is denied in response to an attempt to set an attribute, the ATTRIBUTE bit shall be set to one.

The OBJECT TYPE field shall be set to the object type shown in table 17 (see 4.11.2.2.1) of the highest member of the object accessibility attributes hierarchy that is denying the requested write access.

Changes in the value of any object accessibility attribute that occur after the command, security parameters, and capabilities have been validated shall not affect processing of the command.

4.8 OSD object attributes

4.8.1 Overview

File systems and other systems based on the traditional storage model (see 4.3) store both user data and meta data. OSD object attributes allow the association of meta data with any OSD object (i.e., root, partition, collection, or user). Attributes may be used to describe specific characteristics of an OSD object (e.g., the total amount of

bytes occupied by the OSD object (including attributes), logical size of the OSD object, and the time the OSD object was last modified).

Any OSD command may retrieve attributes and any OSD command may store attributes.

The GET ATTRIBUTES command (see 6.18) and SET ATTRIBUTES command (see 6.36) allow attributes to be retrieved and stored without performing other command functions (see 3.1.10).

The LIST command (see 6.20) and LIST COLLECTION command (see 6.21) allow application clients to retrieve attributes in the attributes pages associated with the objects being listed. Commands that use collections to affect multiple user objects (see 4.6.6.6) allow application clients to retrieve and/or store attributes in the attributes pages associated each object in the collection to which the command is addressed. Otherwise, an OSD command may only retrieve or store attributes in the attributes pages associated with the OSD object addressed by the command.

Attributes are organized in pages for identification and reference. The attributes within a page have similar sources or uses. Within each attributes page, attributes are identified by an attribute number. Each attributes page is associated with one of the following:

- a) The root object;
- b) A partition;
- c) A collection;
- d) A user object; or
- e) Any OSD object type (see table 9 in 4.8.5).

With the exception of attributes pages in the attributes page number range assigned to any OSD object types, the same attributes page shall not be associated with more than one OSD object type.

The structures of attributes pages are described by standards (e.g., this standard, other American National Standards, ISO standards), by OSD applications specifications (e.g., SAN file systems, data base systems, fixed data repositories), by publicly available manufacturer specifications, and by other documentation. A range of vendor specific attributes pages is provided for which the usage is not restricted by this standard.

Attributes for which the OSD logical unit is able to return a value (i.e., attributes with a non-zero length) are called defined attributes. Attributes with no returnable value (i.e., attributes with a zero length) are called undefined attributes.

4.8.2 Attribute length when retrieving defined and undefined attributes

When an attribute length is returned as part of the retrieved attribute information for a defined attribute (see 3.1.14), the attribute length is the non-zero length of the attribute value.

Attempts to retrieve an undefined attribute (see 3.1.51) shall result in an attribute length of zero being returned (see 7.1.4.3 and 7.1.4.4). The number of bytes in the attribute value field is zero and six pad bytes are included to eight-byte align the descriptor.

4.8.3 Attribute length when setting defined and undefined attributes

The combination of the length to which an attribute is being set using CDB set attributes parameters (see 5.2.6) and whether the OSD logical unit knows the attribute as a defined attribute (see 3.1.14) or an undefined attribute (see 3.1.51) affect an attribute setting operation as shown in table 8.

Table 8 — Setting defined attributes and undefined attributes

Attribute length to be set	Defined attribute	Undefined attribute
zero	Change the attribute from being a defined attribute to being an undefined attribute.	Leave the attribute as an undefined attribute. This shall not be considered an error.
non-zero	Replace the current value of the defined attribute with the specified new value. This shall not be considered an error.	Change the attribute from being an undefined attribute to being a defined attribute with the specified value.

If this standard specifies the length of an attribute, any command that attempts to set that attribute's length to a value other than what this standard specifies (e.g., a command that attempts to undefine the user object length attribute described in 7.1.3.11) shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set as follows:

- a) If the invalid attribute length is in a CDB field, the additional sense code shall be set to INVALID FIELD IN CDB; or
- b) If the invalid attribute length is in the Data-Out Buffer, the additional sense code shall be set to INVALID FIELD IN PARAMETER LIST.

4.8.4 Command function ordering for commands that get and/or set attributes

OSD commands provide the application client with the ability to get and set attributes as part of processing the command (e.g., a WRITE command may also retrieve the user object logical length attribute). This subclause defines the relative order of the command functions (see 3.1.10) processing within a single command.

For commands that use collections to affect multiple objects (see 4.6.6.6) the relative order of command functions processing described in this subclause applies to how the command functions are performed on each user object in the collection (i.e., the order described in this subclause is applied fully to the first user object, then the second user object, then the third, etc.).

Commands other than GET ATTRIBUTES, GET MEMBER ATTRIBUTES, SET ATTRIBUTES, SET MEMBER ATTRIBUTES, REMOVE, REMOVE MEMBER OBJECTS, REMOVE PARTITION, and REMOVE COLLECTION that include getting or setting attributes shall be processed in the following order:

- 1) Process those command functions not related to attributes (e.g., writing data to a user object);
- 2) Process any set attributes command functions resulting from the processing of the command (e.g., changes due to a WRITE command);
- 3) Process any set attributes command functions specified in the CDB; and
- 4) Process any get attributes command functions specified in the CDB.

A GET ATTRIBUTES command shall be processed in the following order:

- 1) Process any set attributes command functions resulting from the processing of the command (e.g., updating the attributes related timestamps);
- 2) Process any get attributes command functions specified in the CDB; and
- 3) Process any set attributes command functions specified in the CDB.

A GET MEMBER ATTRIBUTES command shall be processed in the following order:

- 1) Process any set attributes command functions resulting from the processing of the command (e.g., updating the attributes related timestamps);
- 2) Process any get attributes command functions specified in the CDB for the user object members of the collection;
- 3) Process any set attributes command functions specified in the CDB for the user object members of the collection;
- 4) Process any get attributes command functions specified in the CDB for the collection; and
- 5) Process any set attributes command functions specified in the CDB for the collection.

A SET ATTRIBUTES command shall be processed in the following order:

- 1) Process any set attributes command functions resulting from the processing of the command (e.g., updating the attributes related timestamps);
- 2) Process any set attributes command functions specified in the CDB; and
- 3) Process any get attributes command functions specified in the CDB.

A SET MEMBER ATTRIBUTES command shall be processed in the following order:

- 1) Process any set attributes command functions resulting from the processing of the command (e.g., updating the attributes related timestamps);
- 2) Process any set attributes command functions specified in the CDB for the user object members of the collection;
- 3) Process any get attributes command functions specified in the CDB for the user object members of the collection;
- 4) Process any set attributes command functions specified in the CDB for the collection; and
- 5) Process any get attributes command functions specified in the CDB for the collection.

A REMOVE command, a REMOVE MEMBER OBJECTS command, a REMOVE PARTITION command, or a REMOVE COLLECTION command that includes getting or setting attributes shall be processed in the following order:

- 1) Process any set attributes command functions specified in the CDB;
- 2) Process any get attributes command functions specified in the CDB;
- 3) Process those command functions not related to attributes; and
- 4) Process any set attributes command functions resulting from the processing of the command (e.g., updating timestamps).

4.8.5 Attributes pages

Each attributes page contains attributes with similar sources or uses. Identifying numbers are assigned to attributes pages with ranges of page numbers (see table 9) indicating the type of OSD object with which an attributes page is associated.

Table 9 — Attributes page numbers

Page Number	OSD object type with which the attributes page is associated
0h to 2FFF FFFFh	User
3000 0000h to 5FFF FFFFh	Partition
6000 0000h to 8FFF FFFFh	Collection
9000 0000h to BFFF FFFFh	Root
C000 0000h to EFFF FFFFh	Reserved
F000 0000h to FFFF FFEEh	Any OSD object type (i.e., root, partition, collection, or user)
FFFF FFFFh	Any OSD object type ^a
^a Attributes page number FFFF FFFFh is used to request the retrieval of all attributes pages for a given OSD object type.	

For attributes pages associated with partitions, collections, or the root object, the following constant values are used in this standard:

- a) P is equal to 3000 0000h (e.g., P+5h means 3000 0005h);
- b) C is equal to 6000 0000h (e.g., C+3h means 6000 0003h); and
- c) R is equal to 9000 0000h (e.g., R+2h means 9000 0002h).

No constant is needed for attributes pages that are associated with user objects.

Except for the attributes page numbers that apply to any OSD object type (i.e., F000 0000h through FFFF FFFFh), the ranges of attributes page numbers shown in table 9 are subdivided as shown in table 10.

Table 10 — Attributes page number sets

Page Number Within Range	Description
0h to 7Fh	Defined by this standard
80h to 7FFFh	Reserved
8000h to EFFFh	Defined by other standards (see Annex A)
F000h to FFFFh	Defined by publicly available manufacturer specifications
1 0000h to 1FFF FFFFh	Assigned by the OSD logical unit ^a
2000 0000h to 2FFF FFFFh	Vendor specific
^a The attributes in these pages are undefined (see 3.1.51) until they are set (see 4.8.3). The attribute number 0h should be set as specified in 7.1.3.2 to maintain correct attribute directory information. The OSD logical unit shall not modify attribute values in these pages except in response to information provided in the set attributes parameters in a CDB.	

Attributes pages contain attributes (see 4.8.6). For an example of an attributes page containing attributes see 7.1.3.16.

See 7.1.3 for information about attributes pages defined by this standard, including attributes page numbers that apply to any OSD object type.

4.8.6 Attributes

Each attribute within an attributes page (see 4.8.5) has a unique number between 0h and FFFF FFFEh. The description of each attribute defines the format and usage of that attribute. For examples of attribute definitions see 7.1.3.

The attribute with the attribute number 0h contains the name of the page in the format described in 7.1.3.2. The attribute number FFFF FFFFh may be used to request the retrieval of all the defined attributes (see 3.1.14) in a page.

If a command attempts to set attribute number FFFF FFFFh, it shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set as follows:

- a) If the invalid attribute length is in a CDB field, the additional sense code shall be set to INVALID FIELD IN CDB; or
- b) If the invalid attribute length is in the Data-Out Buffer, the additional sense code shall be set to INVALID FIELD IN PARAMETER LIST.

4.8.7 Attributes directories

The root object, partitions, collections, and user objects shall have associated attributes directory pages as described in table 11.

Table 11 — Attributes directory pages

Page Number	Page Name	Attributes Page Contents	Reference
R+0h	Root Directory	Contains one attribute for every attributes page associated with the root object.	7.1.3.4
P+0h	Partition Directory	Contains one attribute for every attributes page associated with the partition.	7.1.3.5
C+0h	Collection Directory	Contains one attribute for every attributes page associated with the collection.	7.1.3.6
0h	User Object Directory	Contains one attribute for every attributes page associated with the user object.	7.1.3.7

Attributes directory pages shall be maintained by the OSD logical unit.

Application clients may modify an attributes directory page by modifying the contents of attribute number 0h in an attributes page other than the attributes directory page. The definitions for attributes pages with page numbers that are not assigned by the OSD logical unit may prohibit changes in attribute number 0h in order to make their directory entries unchangeable. Any command that attempts to modify an attributes directory page in any other manner shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB or INVALID FIELD IN PARAMETER LIST as appropriate.

Attributes pages that are not associated with any object type (i.e., attributes pages with page numbers between F000 0000h and FFFF FFFFh inclusive) do not appear in any attributes directory.

4.9 Command atomicity and isolation

4.9.1 Overview

Atomicity refers to the number of bytes that the OBSD writes to stable storage (see 4.14) as a group in a manner that ensures that all the bytes or none of them are written (e.g., in a traditional block storage device the number of bytes of atomicity is one block or 512 bytes).

Isolation refers to the degree of interaction between concurrent commands.

Atomicity and isolation interact when atomicity byte count affects the degree of isolation (e.g., if the RANGE isolation method described in table 162 (see 7.1.3.8) uses the data atomicity guarantee to identify which commands overlap for isolation purposes).

4.9.2 Atomicity

The atomicity guarantees described in this standard apply to the following (in priority order):

- 1) Commands being processed for which GOOD status has not yet been returned; and
- 2) Commands for which GOOD status has been returned, with the caveat that an application client has no way to distinguish between atomicity effects and effects caused by media failures after it receives the GOOD status.

If a media error is detected while a command is being processed, the atomicity guarantees affect how much data may have been transferred before the error was detected.

If data transfer errors (e.g., media errors, NVRAM failures, power loss) cause a data loss after GOOD status has been returned, then the atomicity guarantees provide one of the boundaries for which data may have been lost. Detecting and recovering from such errors is described in 4.11.3.

Atomicity also affects performance, but control over these effects is outside the scope of this standard.

Application clients use the values in the Root Information attributes page (see 7.1.3.8) attributes shown in table 12 to tailor the commands they send with respect to the atomicity properties of the OBSD.

Table 12 — Atomicity attributes

Attribute	Name	Description
Data atomicity guarantee	D_LIMIT	The minimum number of D_ALIGN aligned user data bytes that the device server shall write to stable storage as a group.
Data atomicity alignment	D_ALIGN	The data alignment value that maximizes the D_LIMIT number of user data bytes that the device server shall write to stable storage as a group. If D_ALIGN is set to zero, then it is processed as if it is set to one.
Attributes atomicity guarantee	A_LIMIT	The minimum number of bytes of an application client settable attribute that the device server shall write to stable storage as a group.
Data/attributes atomicity multiplier	DA_MULT	The multiplier applied to the sum of D_LIMIT and A_LIMIT to determine the minimum number of combined user data and application client settable bytes that the device server shall write to stable storage as a group in response to one command. If either D_LIMIT or A_LIMIT is set to zero, then DA_MULT shall be processed as if it is set to zero.

D_LIMIT and D_ALIGN combine as shown in the following formula to indicate the minimum number of user data bytes that the device server guarantees to write to stable storage as a group:

$$\text{min_bytes} = \text{D_LIMIT} - \text{the_remainder_from}(\text{starting byte offset} \div \text{D_ALIGN})$$

A_LIMIT has no effect on how OSD logical unit provided attributes are written to stable storage. All OSD logical unit provided attributes shall be written to stable storage in a manner that maximizes their integrity and consistency.

Whether the bytes in two or more attributes are written to stable storage as a group is outside the scope of this standard.

If one command writes user data and sets attributes, the D_LIMIT, D_ALIGN, and A_LIMIT attributes still apply, and the minimum total number of bytes that the device server guarantees to write to stable storage as a group is computed as follows:

$$\text{min_tot_bytes} = \text{DA_MULT} \times (\text{D_LIMIT} - \text{the_remainder_from}(\text{starting byte offset} \div \text{D_ALIGN}) + \text{A_LIMIT})$$

Bytes are written to stable storage in multiple groups only when their numbers exceed the guaranteed minimums.

Table 13 shows examples of atomicity attributes effects.

Table 13 — Examples of atomicity attributes effects

Atomicity attribute values				Effect ^a
D_LIMIT	D_ALIGN	A_LIMIT	DA_MULT	
0	1	0	0	The device server provides no guarantees regarding the number of bytes written to stable storage as a group.
512	512	0	0	At least 512 bytes of user data aligned on a 512-byte boundary are written to stable storage as a group.
512	512	512	0	At least 512 bytes of user data aligned on a 512-byte boundary are written to stable storage as a group. At least 512 bytes of an application client settable attribute are written to stable storage as a group. The user data may be written in a separate group from the attribute.
512	512	512	1	At least 512 bytes of user data aligned on a 512-byte boundary at least 512 bytes of an application client settable attribute are written to stable storage in one group.
^a The effects shown in this example all assume that more than 512 bytes of user data and more than 512 bytes of an attribute value are requested to be written.				

4.9.3 Isolation

This standard defines several isolation methods in table 162 (see 7.1.3.8).

The default isolation method attribute in the Root Information attributes page (see 7.1.3.8) specifies the isolation method used by commands that do not override the default. Commands may override the default isolation method by specifying a non-zero value in the ISOLATION field (see 5.2.8).

4.10 Quotas

4.10.1 Introduction

The root, partition, and user objects include attributes pages (see 4.8) that define limits on an application client's ability to consume OSD logical unit resources. The attributes pages are the:

- a) Root Quotas attributes page (see 7.1.3.12);
- b) Partition Quotas attributes page (see 7.1.3.13); and
- c) User Object Quotas attributes page (see 7.1.3.14).

The command and attributes definitions in this standard (see 5.2.6, clause 6, and 7.1.3) specify which quotas are to be tested and how they are to be tested for each command or attribute capable of generating a quota error.

4.10.2 Quota errors

If one of the quota error conditions described in 5.2.6, clause 6, or 7.1.3 occurs, the command shall be terminated with CHECK CONDITION status, with the sense key set to DATA PROTECT and the additional sense code set to QUOTA ERROR. The sense data shall include the OSD attribute identification sense data descriptor (see 4.16.2.3) with one or more attribute descriptors identifying the quota attribute or attributes that have been exceeded.

For multi-object commands, quota error processing shall be handled as described in 4.6.6.6. For individual objects, the device server shall not terminate a command for quota errors after any user data or attributes have been modified.

4.10.3 Quota testing

The device server may implement a vendor specific margin in the tests related to any quota and generate a quota error if a command attempts to consume resources within the margin adjusted quota limit. The size of the vendor specific margin may vary over time in a vendor specific manner.

4.10.4 Changing quotas

The quota values are contained in attributes that may be set by a command with an appropriate capability (see 4.11.2). The device server may constrain the values to which a quota attribute may be set and return CHECK CONDITION status if an attempt is made to set a quota to an unsupported value.

Setting a quota to a value that is less than the applicable resources already consumed in the OSD logical unit:

- a) Shall not be an error; and
- b) Shall not result in the truncation or removal of any information (e.g., user data or attribute values) already stored by the OSD logical unit.

As long as the quota value remains set to a value that is less than the applicable resources already consumed, all commands that attempt to consume additional applicable resource shall be terminated with a quota error.

4.11 Policy/storage management

4.11.1 Overview

The policy/storage manager:

- a) Provides access policy controls to application clients via preparation of policy-coordinated capabilities (see 4.11.2); and
- b) In concert with the OSD logical unit:
 - A) Identifies damaged storage within the OBSD (see 4.11.3.2);
 - B) Repairs damage that the OSD logical unit is unable to repair without assistance (see 4.11.3.1);
 - C) Uses access policy controls (see 4.11.3.2), unit attention conditions (see 4.11.3.1), and CHECK CONDITION status (see 4.11.3.3) to prevent unsafe or temporarily undesirable utilization of OBSD storage.

4.11.2 Capabilities

4.11.2.1 Introduction

Each CDB defined by this standard includes a solo capability (see 5.2.1) whose contents specify the command functions (see 3.1.10) that the device server is allowed to process in response to the command. If allowed by a command's definition, the CDB continuation segment (see 5.3) may be used to add one or more extension capabilities (see 5.4.6) to the command processing inputs.

The device server validates that the requested command functions are allowed by a solo capability or an extension capability based on:

- a) The type of functions (e.g., read, write, attributes setting, attributes retrieval); and
- b) The OSD object or objects on which the command functions are to be processed.

The policies that determine which capabilities are provided to which application clients are outside the scope of this standard.

The policy/storage manager shall coordinate the delivery of capabilities to application clients with the security manager (see 4.12) as follows:

- a) If the security method for all partitions in the OSD logical unit is NOSEC (see 4.12.4.2), then the policy/storage manager may:
 - A) Allow application clients to prepare their own capabilities;
 - B) Coordinate the preparation of capabilities for multiple application clients in response to requests, the format and transport mechanisms for which are outside the scope of this standard; or
 - C) Coordinate the preparation of capabilities with the security manager as described in item b);or
- b) If a security method other than NOSEC is in use by any partition in the OSD logical unit, then the policy/storage manager shall coordinate the preparation of capabilities with the security manager by:
 - A) Requiring application clients to request credentials and capabilities from the security manager; and
 - B) Preparing capabilities only in response to requests from the security manager.

4.11.2.2 Capability format

4.11.2.2.1 Introduction

A capability (see table 14) is included in a CDB to enable the device server to verify that the sender is allowed to perform the command functions (see 3.1.10) described by the CDB.

Table 14 — Capability format

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved				CAPABILITY FORMAT (2h)			
1	KEY VERSION				INTEGRITY CHECK VALUE ALGORITHM			
2	Reserved				SECURITY METHOD			
3	Reserved							
4	(MSB) _____ CAPABILITY EXPIRATION TIME _____ (LSB)							
9								
10	AUDIT _____							
29								
30	(MSB) _____ CAPABILITY DISCRIMINATOR _____ (LSB)							
41								
42	(MSB) _____ OBJECT CREATED TIME _____ (LSB)							
47								
48	OBJECT TYPE							
49								
53	PERMISSIONS BIT MASK _____							
54	Reserved							
55	OBJECT DESCRIPTOR TYPE				Reserved			
56	(MSB) _____ ALLOWED ATTRIBUTES ACCESS _____ (LSB)							
59								
60								
103	OBJECT DESCRIPTOR _____							

The CAPABILITY FORMAT field (see table 15) specifies the format of the capability. If capabilities are coordinated with the security manager, the capability format also is the credential format. The policy/storage manager shall set the CAPABILITY FORMAT field to 2h (i.e., the format defined by this standard).

Table 15 — Capability format values

Value	Description
0h	No capability
1h	Obsolete
2h	The format defined by this standard
3h to Fh	Reserved

If the CAPABILITY FORMAT field contains 2h, the device server shall verify that the command functions requested by a CDB and any CDB continuation descriptors (see 5.3) are permitted by at least one of the capabilities associated with the command (i.e., by the capability in the CDB (see 5.2.1) or by a capability in the extension capabilities CDB continuation descriptor (see 5.4.6)) as described in this subclause. The device server may verify that a command function is permitted after other command functions are completed. The device server shall verify that a command function is permitted before any part of the command function is performed. (E.g., the device server may delay verifying that the set attributes command functions specified by a set attributes list are allowed until the requested read command function is completed, but all the capability permissions concerning the setting attributes are to be verified before any attribute values are changed.)

The KEY VERSION field, INTEGRITY CHECK VALUE ALGORITHM field, and SECURITY METHOD field are used by the security manager (see 4.12.3). If capabilities are not coordinated with the security manager, the KEY VERSION field, INTEGRITY CHECK VALUE ALGORITHM field, and SECURITY METHOD field are reserved.

If CDB contains a non-zero value in the SECURITY METHOD field, the integrity of the CDB and CDB continuation segment, if any, shall be validated (see 4.12.6.1) before any other command processing actions are undertaken (i.e., before verifying that command functions requested in the CDB are permitted by the capability).

The command shall be terminated as described in 4.11.2.2.5, if the CDB SECURITY METHOD field or CAPABILITY FORMAT field contains zero and one of the following is true:

- a) The command is SET KEY (see 6.37) or SET MASTER KEY (see 6.38); or
- b) The default security method attribute in the attributes page that is located as follows based on the contents of the OBJECT TYPE field specifies a default security method other than NOSEC:
 - A) If the capability OBJECT TYPE field contains ROOT, the default security method attribute in the Root Policy/Security attributes page (see 7.1.3.22);
 - B) If the capability OBJECT TYPE field contains PARTITION, the default security method attribute in the Partition Policy/Security attributes page (see 7.1.3.23) for partition zero (see 4.6.4); or
 - C) If the capability OBJECT TYPE field contains COLLECTION or USER, the default security method attribute in the Partition Policy/Security attributes page for the partition whose Partition ID is contained in the capability ALLOWED PARTITION_ID field.

The CAPABILITY EXPIRATION TIME field specifies the value of the clock attribute in the Root Information attributes page (see 7.1.3.8) after which this capability is no longer valid. If a CAPABILITY EXPIRATION TIME field contains a value other than zero and the value of the clock attribute in the Root Information attributes page (see 7.1.3.8) is greater than the value in the CAPABILITY EXPIRATION TIME field, the command shall be terminated as described in 4.11.2.2.5.

Successful use of the capability expiration time requires some degree of synchronization between the clocks of the device server, policy/storage manager, and security manager. The protocol for synchronizing the clocks is outside the scope of this standard.

The AUDIT field is a vendor specific value that the policy/storage manager and/or security manager may use to associate the capability and credential with a specific application client.

The CAPABILITY DISCRIMINATOR field contains a nonce (see 3.1.24) that differentiates one capability and credential from another.

The OBJECT CREATED TIME field specifies the contents of the created time attribute for the OSD object (see table 16) to which the capability applies. A value of zero specifies that any object created time is allowed.

Table 16 — Created time for OSD objects by type

Object Type (see table 17)	Attributes page containing created time attribute to which the capability OBJECT CREATED TIME field is applies
ROOT	Partition Timestamps attributes page (see 7.1.3.16) for partition zero (see 3.1.33)
PARTITION	Partition Timestamps attributes page
COLLECTION	Collection Timestamps attributes page (see 7.1.3.17)
USER	User Object Timestamps attributes page (see 7.1.3.18)

If a CDB OBJECT CREATED TIME field contains a value other than zero and the value in the OBJECT CREATED TIME field is not identical to the value in the created time attribute from the associated timestamps attributes page (see table 16), then the command shall be terminated as described in 4.11.2.2.5.

The OBJECT TYPE field (see table 17) specifies the type of OSD object to which this capability allows access and aids in the determination of how to validate the capability. If capabilities are coordinated with the security manager, the OBJECT TYPE field is used to select the secret key that is used in validating the credential.

Table 17 — Object type values

Value	Name	OSD object type to which access is allowed
01h	ROOT	Root object
02h	PARTITION	Partition
40h	COLLECTION	Collection
80h	USER	User objects
all other values	Reserved	

If the command functions specified by the CDB and CDB continuation segment, if any, are not allowed for the OSD object type specified in the OBJECT TYPE field of any capability associated with the command (i.e., the capability in the CDB (see 5.2.1) and the capabilities, if any, in the CDB continuation segment (see 5.3)), the command shall be terminated as described in 4.11.2.2.5.

The PERMISSIONS BIT MASK field (see table 18) specifies which functions are allowed by this capability. More than one permissions bit may be set within the constraints specified in 4.11.2.3 resulting in a single capability that allows more than one command function.

Table 18 — Permissions bit mask format

Bit Byte	7	6	5	4	3	2	1	0
49	READ	WRITE	GET_ATTR	SET_ATTR	CREATE	REMOVE	OBJ_MGMT	APPEND
50	DEV_MGMT	GLOBAL	POL/SEC	M_OBJECT	QUERY	GBL_REM	Reserved	
51	Reserved							
52	Reserved							
53	Reserved							

A READ bit set to one allows read access to the data in an OSD object, but not to the attributes. For the root object, partitions, and collections the data in the OSD object is the list of other objects contained in the OSD object. A READ bit set to zero prohibits read access to the data in an OSD object.

A WRITE bit set to one allows processing of the WRITE command (see 6.40) or an equivalent command, but not access to user object attributes. A WRITE bit set to zero prohibits processing of the WRITE command or an equivalent command.

A GET_ATTR (get attributes) bit set to one allows retrieval of (i.e., read access to) the attributes associated with an OSD object. A GET_ATTR bit set to zero prohibits retrieval of attributes except for the attributes in the Current Command attributes page (see 7.1.3.31).

A SET_ATTR (set attributes) bit set to one allows the setting of (i.e., write access to) the attributes associated with an OSD object except for attributes located in the OSD object's policy/security attributes page (e.g., the User Object Policy/Security attributes page (see 7.1.3.25) if the OSD object is a user object). The setting of attributes located in the OSD object's policy/security attributes page is allowed only if both the SET_ATTR bit and the POL/SEC bit are set to one. A SET_ATTR bit set to zero prohibits the setting of the attributes associated with an OSD object.

A CREATE bit set to one allows the creation of OSD objects. A CREATE bit set to zero prohibits the creation of OSD objects.

A REMOVE bit set to one allows the removal of OSD objects. A REMOVE bit set to zero prohibits the removal of OSD objects.

An OBJ_MGMT (object management) bit set to one allows command functions that may change how the OSD logical unit handles an OSD object without affecting the stored data, stored attributes, commands in the task set, policies, or security for the OSD object. An OBJ_MGMT bit set to zero prohibits such command functions.

An APPEND bit set to one allows processing of the APPEND command (see 6.2), but not access to user object attributes. A APPEND bit set to zero prohibits processing of the APPEND command.

A DEV_MGMT (device management) bit set to one allows command functions that affect the OSD logical unit. A DEV_MGMT bit set to zero prohibits command functions that affect the OSD logical unit.

A GLOBAL bit set to one allows command functions that may affect all the OSD objects in the OSD logical unit. A GLOBAL bit set to zero prohibits command functions that may affect all the OSD objects in the OSD logical unit.

A POL/SEC bit set to one allows command functions that affect the policy/security functions performed for one or more OSD objects. A POL/SEC bit set to zero prohibits command functions that affect the policy/security functions performed for one or more OSD objects.

A multiple objects (M_OBJECT) bit set to one in combination with other permissions bits allows retrieving attributes from multiple user objects, setting attributes in multiple user objects, and removing multiple user objects. An M_OBJECT bit set to zero prohibits multiple user object commands.

A QUERY bit set to one allows searching the user objects in a collection for specified attribute values. An QUERY bit set to zero prohibits searching the user objects in a collection.

A GBL_REM (global remove) bit set to one allows all the user objects, collections, and partitions referenced by a single command to be removed (e.g., the global removal of user objects and collections performed by a REMOVE PARTITION command (see 6.34) with the REMOVE SCOPE field set to 001b). A GBL_REM bit set to zero prohibits the removal of all the user objects, collections, and partitions referenced by a single command.

The OBJECT DESCRIPTOR TYPE field (see table 19) specifies the format of information that appears in the OBJECT DESCRIPTOR field.

Table 19 — Object descriptor types

Object Descriptor Type	Name	Description	Reference
0h	NONE	The OBJECT DESCRIPTOR field shall be ignored	
1h	USER	A single user object	4.11.2.2.2
2h	PAR	A single partition, including partition zero	4.11.2.2.3
3h	COL	A single collection	4.11.2.2.4
3h to Fh		Reserved	

The ALLOWED ATTRIBUTES ACCESS field (see table 20) places additional restrictions on the attributes that the command is able to access.

Table 20 — ALLOWED ATTRIBUTES ACCESS field

Value	Description
0h	No additional restrictions are placed on attributes accesses.
1h to FFFF FFEh	The contents of the Attributes Access attributes page attribute for the partition specified by the ALLOWED PARTITION_ID field in the capability object descriptor specified by the ALLOWED ATTRIBUTES ACCESS field restrict the attributes to which access is allowed as described in 7.1.3.19.
FFFF FFFFh	Reserved

If the ALLOWED ATTRIBUTES ACCESS field specifies the attribute number of an attribute that is undefined (see 3.1.51) in the Attributes Access attributes page (see 7.1.3.19) attribute for the partition specified by the ALLOWED PARTITION_ID field in the capability object descriptor, then the command shall be terminated as described in 4.11.2.2.5.

4.11.2.2.2 USER capability object descriptor

If the object descriptor type is USER (i.e., 1h), the OBJECT DESCRIPTOR field shall have the format shown in table 21, specifying a single user object and a range of bytes within that user object to which the capability allows access.

Table 21 — User object descriptor format

Bit Byte	7	6	5	4	3	2	1	0
60	(MSB)	POLICY ACCESS TAG						(LSB)
63								
64	(MSB)	BOOT EPOCH						(LSB)
65								
66		Reserved						
71								
72	(MSB)	ALLOWED PARTITION_ID						(LSB)
79								
80	(MSB)	ALLOWED USER_OBJECT_ID						(LSB)
87								
88	(MSB)	ALLOWED RANGE LENGTH						(LSB)
95								
96	(MSB)	ALLOWED RANGE STARTING BYTE ADDRESS						(LSB)
103								

If the POLICY ACCESS TAG field contains a value other than zero, the policy access tag attribute identified by the command and OBJECT TYPE field (see table 22) is compared to the POLICY ACCESS TAG field contents as part of verifying the capability. If the POLICY ACCESS TAG field contains zero, then no comparison is made to any policy access tag attribute. The policy/storage manager or OSD logical unit changes the policy access tag to prevent unsafe or temporarily undesirable accesses to an OSD object (see 4.11.3.2).

Table 22 — Policy access tag usage for OSD object types and commands

Command	Object Type (see table 17)	Attributes page containing policy access tag attribute to which CDB POLICY ACCESS TAG field is compared
CREATE PARTITION or REMOVE PARTITION	PARTITION	Partition Policy/Security attributes page (see 7.1.3.23) for partition zero (see 3.1.33)
CREATE COLLECTION or REMOVE COLLECTION	COLLECTION	Partition Policy/Security attributes page
CREATE, CREATE AND WRITE, or REMOVE	USER	Partition Policy/Security attributes page
All other commands	ROOT	Partition Policy/Security attributes page for partition zero
	PARTITION	Partition Policy/Security attributes page
	COLLECTION	Collection Policy/Security attributes page (see 7.1.3.24)
	USER	User Object Policy/Security attributes page (see 7.1.3.25)

If the non-zero value in the POLICY ACCESS TAG field is not identical to the value in the policy access tag attribute from the associated policy/security attributes page (see table 22), then the command shall be terminated as described in 4.11.2.2.5.

If the BOOT EPOCH field contains zero or the boot epoch attribute in the Root Policy/Security attributes page (see 7.1.3.22) contains zero, then the contents of the BOOT EPOCH field shall be ignored. If the non-zero values in the BOOT EPOCH field and the boot epoch attribute in the Root Policy/Security attributes page do not match, then the command shall be terminated as described in 4.11.2.2.5.

The ALLOWED PARTITION_ID field specifies the Partition_ID (see 4.6.4) of the partition to which access is allowed. If the ALLOWED PARTITION_ID field contains zero, the command shall be terminated as described in 4.11.2.2.5.

The ALLOWED USER_OBJECT_ID field specifies the User_Object_ID (see 4.6.5) of the OSD object to which the capability allows access. If the ALLOWED USER_OBJECT_ID field contains zero and the command is not CREATE (see 6.5) or CREATE AND WRITE (see 6.6), then the command shall be terminated as described in 4.11.2.2.5.

The ALLOWED RANGE LENGTH field specifies number of bytes in the range of user object bytes to which the capability allows access.

The ALLOWED RANGE STARTING BYTE OFFSET field specifies the location of the first byte in the range of user object bytes to which the capability allows access relative to the first byte (i.e., byte zero).

If the ALLOWED RANGE LENGTH field is set to FFFF FFFF FFFF FFFFh and the ALLOWED RANGE STARTING BYTE OFFSET field is set to zero, then access is allowed to all bytes in the user object.

If the ALLOWED RANGE LENGTH field is set to FFFF FFFF FFFF FFFFh and the ALLOWED RANGE STARTING BYTE OFFSET field is set to a non-zero value, then access is allowed to all bytes from the allowed range starting byte to byte FFFF FFFF FFFF FFFFh. This shall not be considered an error.

The command that accesses a user object shall be terminated as described in 4.11.2.2.5, if none of the capabilities associated with the command (i.e., the capability in the CDB (see 5.2.1) and the capabilities, if any, in the CDB continuation segment (see 5.3)) match all of the following criteria:

- a) The partition that contains the user object (e.g., the partition specified by the PARTITION_ID field in the CDB of a READ command) matches the ALLOWED PARTITION_ID field in the capability;
- b) The user object being accessed (e.g., the user object specified by the USER_OBJECT_ID field in the CDB of a READ command) matches the ALLOWED USER_OBJECT_ID field in the capability; and
- c) If data is allowed to be transferred to or from the user object (i.e., if the READ permission bit or the WRITE permission bit (see 4.11.2.2.1) is set to one), then the specified range of bytes being transferred (e.g., the range of bytes specified by the LENGTH field and STARTING BYTE ADDRESS field in the CDB of a READ command with the CDB CONTINUATION LENGTH field (see 5.2.1) is set to zero) is inside the range of bytes specified by the ALLOWED RANGE LENGTH field and ALLOWED RANGE STARTING BYTE OFFSET field.

4.11.2.2.3 PAR capability object descriptor

If the object descriptor type is PAR (i.e., 2h), the OBJECT DESCRIPTOR field shall have the format shown in table 23, specifying a single partition to which the capability allows access. For a LIST COLLECTION command with the M_OBJECT bit set to one (see 4.11.2.2.1), the PAR capability object descriptor allows access to a single partition and the attributes associated with each collection in the partition. For the LIST command with the M_OBJECT bit set to one, the PAR capability object descriptor allows access to:

- a) The root object and the attributes associated with each partition; or
- b) A partition and the attributes associated with each user object in the partition.

Table 23 — Partition descriptor format

Bit Byte	7	6	5	4	3	2	1	0
60	(MSB)	POLICY ACCESS TAG						(LSB)
63								
64	(MSB)	BOOT EPOCH						(LSB)
65								
66		Reserved						
71								
72	(MSB)	ALLOWED PARTITION_ID						(LSB)
79								
80		Reserved						
103								

The POLICY ACCESS TAG field and BOOT EPOCH field are described in 4.11.2.2.2.

The ALLOWED PARTITION_ID field specifies the partition to which access is allowed. The command shall be terminated as described in 4.11.2.2.5, if any of the following are true:

- a) If the OBJECT TYPE field contains 02h (i.e., PARTITION), the ALLOWED PARTITION_ID field contains zero, and the command is not one of the following:
 - A) CREATE PARTITION (see 6.9);
 - B) CREATE CLONE (see 6.7); or
 - C) CREATE SNAPSHOT (see 6.10);
 or
- b) If the OBJECT TYPE field contains 01h (i.e., ROOT) and the ALLOWED PARTITION_ID field contains a value other than zero.

The command that accesses a partition or a well known collection (see 4.6.6.5) in a partition shall be terminated as described in 4.11.2.2.5, if none of the capabilities associated with the command (i.e., the capability in the CDB (see 5.2.1) and the capabilities, if any, in the CDB continuation segment (see 5.3)) match all of the following criteria:

- a) If the OBJECT TYPE field contains:
 - A) 02h (i.e., PARTITION), then the partition being accessed (e.g., the partition specified by the PARTITION_ID field in the CDB of a LIST command) matches the ALLOWED PARTITION_ID field in the capability; or

- B) 01h (i.e., ROOT), then the partition being accessed (e.g., the partition specified by the PARTITION_ID field in the CDB of a LIST command) is zero;
- and
- b) The User_Object_ID (see 4.6.2) associated with the object being accessed, if any, is:
 - A) Zero; or
 - B) The Collection_Object_ID of a well known collection (see 4.6.6.5).

4.11.2.2.4 COL capability object descriptor

If the object descriptor type is COL (i.e., 3h), the OBJECT_DESCRIPTOR field shall have the format shown in table 24, specifying a single collection to which the capability allows access. If the M_OBJECT permission bit is set to one or the QUERY permission bit is set to one (see 4.11.2.2.1), the COL capability object descriptor allows access to a single collection and the attributes associated with each user object in the collection.

Table 24 — Collection descriptor format

Bit Byte	7	6	5	4	3	2	1	0
60	(MSB)							
63	POLICY ACCESS TAG							(LSB)
64	(MSB)							
65	BOOT EPOCH							(LSB)
66	Reserved							
71								
72	(MSB)							
79	ALLOWED PARTITION_ID							(LSB)
80	(MSB)							
87	ALLOWED COLLECTION_OBJECT_ID							(LSB)
88	Reserved							
103								

The POLICY ACCESS TAG field, BOOT EPOCH field, and ALLOWED PARTITION_ID field are described in 4.11.2.2.2.

The ALLOWED COLLECTION_OBJECT_ID field specifies the Collection_Object_ID (see 4.6.6) of the collection to which the capability allows access. If the ALLOWED COLLECTION_OBJECT_ID field contains zero and the command is not CREATE COLLECTION (see 6.8), then the command shall be terminated as described in 4.11.2.2.5.

The command that accesses a collection shall be terminated as described in 4.11.2.2.5, if none of the capabilities associated with the command (i.e., the capability in the CDB (see 5.2.1) and the capabilities, if any, in the CDB continuation segment (see 5.3)) match all of the following criteria:

- a) The partition that contains the collection (e.g., the partition specified by the PARTITION_ID field in the CDB of a SET MEMBER ATTRIBUTES command) matches the ALLOWED PARTITION_ID field in the capability; and
- b) The collection being accessed (e.g., the collection specified by the COLLECTION_OBJECT_ID field in the CDB of a SET MEMBER ATTRIBUTES command) matches the ALLOWED COLLECTION_OBJECT_ID field in the capability.

4.11.2.2.5 Command termination due to errors detected in a capability

If an error is detected during the validation of a capability, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set as follows:

- a) If the capability in which the error is detected is in the CDB (see 5.2.1), the additional sense code shall be set to INVALID FIELD IN CDB; or
- b) If the capability in which the error is detected is in the CDB continuation segment (see 5.3), the additional sense code shall be set to INVALID FIELD IN PARAMETER LIST.

4.11.2.3 Capabilities and commands allowed

The validity of a specific command and some of the command function (see 3.1.10) related fields in that command is determined by the presence of specific combinations of values in capability fields as shown in table 25. A command function is allowed if at least one row in table 25 allows it, even if a different row that applies does not allow it.

Any command may retrieve or set attributes. The combinations of capability fields that allow those functions are shown in table 26. Retrieving or setting attributes is allowed if at least one row in table 26 allows it, even if a different row that applies does not allow it.

A single capability for a single object type may allow processing of multiple command functions (e.g., read and write) as well as the retrieving and setting of attributes by combining the permission bits values described in multiple rows of table 25 and table 26.

Table 25 — Commands allowed by specific capability field values (part 1 of 5)

Commands allowed and CDB fields whose contents are restricted by capability field contents, if any	Capability Field values that allow a command		
	Object Type Name	Permission Bits That Are Set To One	Object Descriptor Name
An APPEND command	USER	APPEND	USER
A CLEAR command	USER	WRITE	USER
A COPY USER OBJECTS command with one destination user object and one or more source user objects ^a			
Destination user object	USER	CREATE and WRITE	USER
Source user object or user objects	USER	READ	USER
A CREATE command	USER	CREATE	USER
Combinations of OBJECT TYPE field, PERMISSION BITS field, and OBJECT DESCRIPTOR TYPE field values not shown in this table and table 26 are reserved. The capability fields not shown in this table may place additional limits on the objects that are allowed to be accessed.			
^a This command accesses multiple objects. One capability is necessary for each object accessed. The solo capability (see 3.1.46) appears in the CDB (see 5.2.1). The other capabilities appear in the CDB continuation segment (see 5.3).			

Table 25 — Commands allowed by specific capability field values (part 2 of 5)

Commands allowed and CDB fields whose contents are restricted by capability field contents, if any	Capability Field values that allow a command		
	Object Type Name	Permission Bits That Are Set To One	Object Descriptor Name
A CREATE AND WRITE command	USER	CREATE and WRITE	USER
A CREATE COLLECTION command	COLLECTION	CREATE	COL
A CREATE CLONE command ^a			
Source partition	PARTITION	READ	PAR
Destination partition	PARTITION	WRITE	PAR
A CREATE PARTITION command	PARTITION	CREATE	PAR
A CREATE SNAPSHOT command ^a			
Source partition	PARTITION	READ	PAR
Destination partition	PARTITION	WRITE	PAR
A CREATE USER TRACKING COLLECTION command with the SOURCE COLLECTION_OBJECT_ID field set to zero	COLLECTION	CREATE	COL
A CREATE USER TRACKING COLLECTION command with the SOURCE COLLECTION_OBJECT_ID field set to a non-zero value ^a			
The collection specified by the COLLECTION_OBJECT_ID field in the CDB	COLLECTION	CREATE and WRITE	COL
The collection specified by the SOURCE COLLECTION_OBJECT_ID field in the CDB	COLLECTION	READ	COL
A DETACH CLONE command	PARTITION	WRITE	PAR
A FLUSH command	USER	OBJ_MGMT	USER
A FLUSH COLLECTION command	COLLECTION	OBJ_MGMT	COL
A FLUSH PARTITION command	PARTITION	OBJ_MGMT	PAR
A FLUSH OSD command	ROOT	OBJ_MGMT	PAR
A FORMAT OSD command	ROOT	OBJ_MGMT and GLOBAL	PAR
A GET ATTRIBUTES command addressed to a user object	USER	see table 26	USER
Combinations of OBJECT TYPE field, PERMISSION BITS field, and OBJECT DESCRIPTOR TYPE field values not shown in this table and table 26 are reserved. The capability fields not shown in this table may place additional limits on the objects that are allowed to be accessed.			
^a This command accesses multiple objects. One capability is necessary for each object accessed. The solo capability (see 3.1.46) appears in the CDB (see 5.2.1). The other capabilities appear in the CDB continuation segment (see 5.3).			

Table 25 — Commands allowed by specific capability field values (part 3 of 5)

Commands allowed and CDB fields whose contents are restricted by capability field contents, if any	Capability Field values that allow a command		
	Object Type Name	Permission Bits That Are Set To One	Object Descriptor Name
A GET ATTRIBUTES command addressed to a collection	COLLECTION	see table 26	COL
A GET ATTRIBUTES command addressed to a partition	PARTITION	see table 26	PAR
A GET ATTRIBUTES command addressed to the root object	ROOT	see table 26	PAR
A GET MEMBER ATTRIBUTES command addressed to a collection	COLLECTION	see table 26	COL
A LIST command addressed to a partition with the LIST_ATTR bit to be set to zero	PARTITION	READ	PAR
A LIST command addressed to a partition	PARTITION	READ and M_OBJECT	PAR
A LIST command addressed to the root object	ROOT	READ and M_OBJECT	PAR
A LIST COLLECTION command addressed to a collection with the LIST_ATTR bit to be set to zero	COLLECTION	READ	COL
A LIST COLLECTION command addressed to a collection	COLLECTION	READ and M_OBJECT	COL
A LIST COLLECTION command addressed to a partition with the LIST_ATTR bit to be set to zero	PARTITION	READ	PAR
A LIST COLLECTION command addressed to a partition	PARTITION	READ and M_OBJECT	PAR
A LIST COLLECTION command addressed to a well known collection with the LIST_ATTR bit set to zero	PARTITION	READ	PAR
An OBJECT STRUCTURE CHECK command addressed to a partition	PARTITION	DEV_MGMT	PAR
An OBJECT STRUCTURE CHECK command addressed to the root object	ROOT	DEV_MGMT	PAR
Combinations of OBJECT TYPE field, PERMISSION BITS field, and OBJECT DESCRIPTOR TYPE field values not shown in this table and table 26 are reserved. The capability fields not shown in this table may place additional limits on the objects that are allowed to be accessed.			
^a This command accesses multiple objects. One capability is necessary for each object accessed. The solo capability (see 3.1.46) appears in the CDB (see 5.2.1). The other capabilities appear in the CDB continuation segment (see 5.3).			

Table 25 — Commands allowed by specific capability field values (part 4 of 5)

Commands allowed and CDB fields whose contents are restricted by capability field contents, if any	Capability Field values that allow a command		
	Object Type Name	Permission Bits That Are Set To One	Object Descriptor Name
A PERFORM TASK MANAGEMENT command with function code of ABORT TASK or QUERY TASK addressed to a user object	USER	DEV_MGMT	USER
A PERFORM TASK MANAGEMENT command with function code of ABORT TASK or QUERY TASK addressed to a collection	COLLECTION	DEV_MGMT	COL
A PERFORM TASK MANAGEMENT command with function code of ABORT TASK or QUERY TASK addressed to a partition	PARTITION	DEV_MGMT	PAR
A PERFORM TASK MANAGEMENT command with function code of ABORT TASK or QUERY TASK addressed to the root object	ROOT	DEV_MGMT	PAR
A PERFORM TASK MANAGEMENT command or a PERFORM SCSI COMMAND command.	ROOT	DEV_MGMT and GLOBAL	PAR
A PUNCH command	USER	WRITE	USER
A QUERY command with the MATCHES COLLECTION_OBJECT_ID field set to zero	COLLECTION	QUERY	COL
A QUERY command with the MATCHES COLLECTION_OBJECT_ID field set to a non-zero value ^a			
The collection specified by the COLLECTION_OBJECT_ID field in the CDB	COLLECTION	QUERY	COL
The collection specified by the MATCHES COLLECTION_OBJECT_ID field in the CDB	COLLECTION	WRITE	COL
A READ command	USER	READ	USER
A READ MAP command	USER	DEV_MGMT	USER
A READ MAPS AND COMPARE command ^a			
Each user object that participates in the comparison	USER	DEV_MGMT	USER
A REFRESH SNAPSHOT OR CLONE command ^a			
Source partition	PARTITION	READ	PAR
Destination partition	PARTITION	APPEND	PAR
A REMOVE command	USER	REMOVE	USER
Combinations of OBJECT TYPE field, PERMISSION BITS field, and OBJECT DESCRIPTOR TYPE field values not shown in this table and table 26 are reserved. The capability fields not shown in this table may place additional limits on the objects that are allowed to be accessed.			
^a This command accesses multiple objects. One capability is necessary for each object accessed. The solo capability (see 3.1.46) appears in the CDB (see 5.2.1). The other capabilities appear in the CDB continuation segment (see 5.3).			

Table 25 — Commands allowed by specific capability field values (part 5 of 5)

Commands allowed and CDB fields whose contents are restricted by capability field contents, if any	Capability Field values that allow a command		
	Object Type Name	Permission Bits That Are Set To One	Object Descriptor Name
A REMOVE COLLECTION	COLLECTION	REMOVE	COL
A REMOVE MEMBER OBJECTS command addressed to a collection	COLLECTION	REMOVE and M_OBJECT	COL
A REMOVE PARTITION command with the REMOVE SCOPE field set to 000b	PARTITION	REMOVE	PAR
A REMOVE PARTITION command with the REMOVE SCOPE field set to 001b	PARTITION	REMOVE and GBL_REM	PAR
A RESTORE PARTITION FROM SNAPSHOT command ^a			
Main partition	PARTITION	READ	PAR
Snapshot partition	PARTITION	WRITE	PAR
A SET ATTRIBUTES command addressed to a user object	USER	see table 26	USER
A SET ATTRIBUTES command addressed to a collection	COLLECTION	see table 26	COL
A SET ATTRIBUTES command addressed to a partition	PARTITION	see table 26	PAR
A SET ATTRIBUTES command addressed to the root object	ROOT	see table 26	PAR
A SET MEMBER ATTRIBUTES command addressed to a collection	COLLECTION	see table 26	COL
A SET KEY command with KEY TO SET field equal to 10b or 11b	PARTITION	DEV_MGMT and POL/SEC	PAR
Any SET KEY command with KEY TO SET field equal to 01b	ROOT	DEV_MGMT and POL/SEC	PAR
Any SET MASTER KEY command.	ROOT	DEV_MGMT, POL/SEC, and GLOBAL	PAR
A WRITE command	USER	WRITE	USER
Combinations of OBJECT TYPE field, PERMISSION BITS field, and OBJECT DESCRIPTOR TYPE field values not shown in this table and table 26 are reserved. The capability fields not shown in this table may place additional limits on the objects that are allowed to be accessed.			
^a This command accesses multiple objects. One capability is necessary for each object accessed. The solo capability (see 3.1.46) appears in the CDB (see 5.2.1). The other capabilities appear in the CDB continuation segment (see 5.3).			

Table 26 — Attribute retrieving and setting function allowed by specific capability field values (part 1 of 3)

Attribute-Related Functions Allowed	Capability Field values that allow attribute-related functions		
	Object Type Name	Permission Bits That Are Set To One	Object Descriptor Name
Retrieval of attributes from the Current Command attributes page (see 7.1.3.31)	USER	GET_ATTR	USER
Retrieval of attributes from the Current Command attributes page (see 7.1.3.31)	COLLECTION	GET_ATTR	COL
Retrieval of attributes from the Current Command attributes page	PARTITION or ROOT	GET_ATTR	PAR
Retrieval of attributes from an attributes page associated with a well known collection (see 4.6.6.5)	PARTITION	GET_ATTR	PAR
Retrieval of attributes from any attributes page associated with the addressed user object	USER	GET_ATTR	USER
As part of a CREATE command or CREATE AND WRITE command, the retrieval of attributes from any attributes page associated with any user object created by the command	USER	GET_ATTR	USER
Retrieval of attributes from any attributes page associated with the addressed collection	COLLECTION	GET_ATTR	COL
As part of a CREATE COLLECTION command, the retrieval of attributes from any attributes page associated with the collection	COLLECTION	GET_ATTR	COL
As part of a GET MEMBER ATTRIBUTES command, QUERY command, REMOVE MEMBER OBJECTS command, or SET MEMBER ATTRIBUTES command, the retrieval of attributes from each user object in a collection	COLLECTION	GET_ATTR and M_OBJECT	COL
As part of a LIST COLLECTION command with the LIST_ATTR bit to be set to one, the return in the parameter data of attributes from any attributes page associated with each user object in the collection	COLLECTION	GET_ATTR and M_OBJECT	COL
Retrieval of attributes from any attributes page associated with the addressed partition	PARTITION	GET_ATTR	PAR
As part of a CREATE PARTITION command, the retrieval of attributes from any attributes page associated with the created partition	PARTITION	GET_ATTR	PAR
As part of a LIST COLLECTION command with the LIST_ATTR bit to be set to one, the return in the parameter data of attributes from any attributes page associated with each collection in the partition	PARTITION	GET_ATTR and M_OBJECT	PAR
Combinations of OBJECT TYPE field, PERMISSION BITS field, and OBJECT DESCRIPTOR TYPE field values not shown in this table and table 25 are reserved. The capability fields not shown in this table may place additional limits on the objects that are allowed to be accessed.			

Table 26 — Attribute retrieving and setting function allowed by specific capability field values (part 2 of 3)

Attribute-Related Functions Allowed	Capability Field values that allow attribute-related functions		
	Object Type Name	Permission Bits That Are Set To One	Object Descriptor Name
As part of a LIST command with the LIST_ATTR bit to be set to one, the return in the parameter data of attributes from any attributes page associated with each user object in the partition	PARTITION	GET_ATTR and M_OBJECT	PAR
Retrieval of attributes from any attributes page associated with the root object or in any attributes page associated with partition zero (see 3.1.33)	ROOT	GET_ATTR	PAR
Setting attributes in any attributes page associated with the addressed user object, except attributes in a User Object Policy/Security attributes page (see 7.1.3.25)	USER	SET_ATTR	USER
As part of a CREATE command or CREATE AND WRITE command, the setting of attributes in any attributes page associated with any user object created by the command, except attributes in a User Object Policy/Security attributes page	USER	SET_ATTR	USER
Setting attributes in any attributes page associated with the addressed collection, except attributes in a Collection Policy/Security attributes page (see 7.1.3.24)	COLLECTION	SET_ATTR	COL
As part of a CREATE COLLECTION command, the setting of attributes in any attributes page associated with the collection created by the command, except attributes in the Collection Policy/Security attributes page	COLLECTION	SET_ATTR	COL
As part of a GET MEMBER ATTRIBUTES command or SET MEMBER ATTRIBUTES command, the setting of attributes from each user object in the collection, except attributes in a User Object Policy/Security attributes page (see 7.1.3.25)	COLLECTION	SET_ATTR and M_OBJECT	COL
Setting attributes in any attributes page associated with the addressed partition, except attributes in a Partition Policy/Security attributes page (see 7.1.3.23)	PARTITION	SET_ATTR	PAR
As part of a CREATE PARTITION command, the setting of attributes in any attributes page associated with the partition created by the command, except attributes in the Partition Policy/Security attributes page	PARTITION	SET_ATTR	PAR
Setting attributes in any attributes page associated with the root object, except attributes in a Root Policy/Security attributes page, or setting attributes in any attributes page associated with partition zero, except attributes in a Partition Policy/Security attributes page	ROOT	SET_ATTR	PAR
Combinations of OBJECT TYPE field, PERMISSION BITS field, and OBJECT DESCRIPTOR TYPE field values not shown in this table and table 25 are reserved. The capability fields not shown in this table may place additional limits on the objects that are allowed to be accessed.			

Table 26 — Attribute retrieving and setting function allowed by specific capability field values (part 3 of 3)

Attribute-Related Functions Allowed	Capability Field values that allow attribute-related functions		
	Object Type Name	Permission Bits That Are Set To One	Object Descriptor Name
Setting attributes in any attributes page associated with the addressed user object	USER	SET_ATTR and POL/SEC	USER
As part of a CREATE command or CREATE AND WRITE command, the setting of attributes in any attributes page associated with any user object created by the command	USER	SET_ATTR and POL/SEC	USER
Setting attributes in any attributes page associated with the addressed collection	COLLECTION	SET_ATTR and POL/SEC	COL
As part of a CREATE COLLECTION command, the setting of attributes in any attributes page associated with the collection created by the command	COLLECTION	SET_ATTR and POL/SEC	COL
As part of a GET MEMBER ATTRIBUTES command or SET MEMBER ATTRIBUTES command, the setting of attributes from any user object in the collection	COLLECTION	SET_ATTR, M_OBJECT, and POL/SEC	COL
Setting attributes in any attributes page associated with the addressed partition	PARTITION	SET_ATTR and POL/SEC	PAR
As part of a CREATE PARTITION command, the setting of attributes in any attributes page associated with the partition created by the command	PARTITION	SET_ATTR and POL/SEC	PAR
Setting attributes in any attributes page associated with the root object or setting attributes in any attributes page associated with partition zero	ROOT	SET_ATTR and POL/SEC	PAR
Combinations of OBJECT TYPE field, PERMISSION BITS field, and OBJECT DESCRIPTOR TYPE field values not shown in this table and table 25 are reserved. The capability fields not shown in this table may place additional limits on the objects that are allowed to be accessed.			

4.11.3 OBSD storage damage detection, repair, and undesirable utilization prevention

4.11.3.1 Normal usage storage damage detection and repair

The OBSD device server detects damaged storage when:

- a) An application client initiated operation detects an uncorrectable error; or
- b) A background operation outside the scope of this standard detects an uncorrectable error.

When a device server detects uncorrectable storage damage, it does the following:

- a) Sets the FENCE bit to one as described in 4.11.3.2 in the affected objects;
- b) Summarizes the damage by updating the attributes in the Error Recovery attributes pages of the affected objects (e.g., the User Object Error Recovery attributes page (see 7.1.3.29) is updated if the OSD object is a user object); and
- c) Establishes a unit attention condition (see SAM-4) for the initiator port associated with every I_T nexus as follows:
 - A) If the storage damage affects some, but not all, partitions, then a unit attention condition shall be established for each affected partition with the:
 - a) The sense key set to UNIT ATTENTION;
 - b) The additional sense code set to ERROR RECOVERY ATTRIBUTES HAVE CHANGED; and
 - c) The INFORMATION field set to the Partition_ID of an affected partition;or
 - B) If the storage damage affects the root object or all partitions, then a unit attention condition shall be established affected partition with the:
 - a) The sense key set to UNIT ATTENTION;
 - b) The additional sense code set to ERROR RECOVERY ATTRIBUTES HAVE CHANGED; and
 - c) The INFORMATION field set to zero.

An application client that receives notification of uncorrectable damaged storage should forward the notification to the policy/storage manager.

A policy/storage manager that receives notification of uncorrectable damaged storage should:

- a) Use any information received with the notification and appropriate commands (e.g., the QUERY command (see 6.26), the GET ATTRIBUTES command (see 6.18), the READ MAP command (see 6.28)) to identify appropriate repair actions;
- b) Perform the identified repair actions (e.g., rewrite corrupted data using the WRITE command (see 6.40));
- c) Update the affected Error Recovery attributes page or pages; and
- d) Set the affected FENCE bits to zero (see 4.11.3.2).

Application clients and policy/storage managers also may detect data errors that are invisible to the device server. Utilization of such data may be prevented by changing the VERSION field in the policy access tag attributes in the Policy/Security attributes pages associated with the affected objects (e.g., the User Object Policy/Security attributes page (see 7.1.3.25) if the OSD object is a user object) as described in 4.11.3.2.

During normal operation, the policy/storage manager may use the OBJECT STRUCTURE CHECK command (see 6.22) to validate the OBSD storage structures for a partition or the root object (i.e., the entire OBSD), however, the object structures being validated by the OBJECT STRUCTURE CHECK command are inaccessible as described in 6.22 for the duration of command processing.

The device server shall not require the application client to send an OBJECT STRUCTURE CHECK command except as described in 4.11.3.3.

4.11.3.2 Policy access tags

The policy access tag (see table 27) allows the coordinated actions of both the OSD logical unit and policy/storage manager to prevent unsafe or temporarily undesirable utilization of OBSD storage that is assigned to the OSD logical unit.

Table 27 — Policy access tag format

Bit Byte	7	6	5	4	3	2	1	0
0	FENCE	(MSB)						
1								
2								
3								(LSB)

During normal operation the value of the FENCE bit is zero.

If the OSD logical unit detects a condition that would make further accesses to one or more OSD objects unsafe, it shall set the FENCE bit to one in the policy access tag attributes in the Policy/Security attributes pages associated with those objects (e.g., the User Object Policy/Security attributes page (see 7.1.3.25) if the OSD object is a user object) and notify the policy/storage manager of a condition needing attention (see 4.11.3.1). The OSD logical unit, policy/storage manager, or both act to correct whatever conditions are making accesses to the OSD objects unsafe. After the conditions making accesses to the OSD objects unsafe are corrected the policy/storage manager sets the FENCE bit to zero.

If a command contains a request to set the FENCE bit to one, then the command shall be terminated as described in 7.1.2.

To block capability-based access to one or more OSD objects, the policy/storage manager changes the VERSION field in the policy access tag attributes in the Policy/Security attributes pages associated with those objects. The conditions under which the policy/storage manager may be called on to do this include:

- a) Recovery from errors other than those detected by the OSD logical unit that make accesses to one or more OSD object unsafe; and
- b) Receipt of a request to change the policy access tag from the security manager (see 4.12.6.4).

If a command contains a request to set the VERSION field to zero, then the command shall be terminated as described in 7.1.2.

The OSD logical unit shall not modify the contents of a policy access tag VERSION field.

The device server terminates any command received with a capability (see 3.1.4) whose POLICY ACCESS TAG field contains a non-zero value that differs from the policy access tag attribute value in the Policy/Security attributes page associated with the object (see 4.11.2.2).

4.11.3.3 Storage damage detection and repair after a reset

After a hard reset SCSI device condition established in response to an event (see SAM-4), the device server may oblige the application client to send an OBJECT STRUCTURE CHECK command (see 6.22) for:

- a) One or more individual partitions; or
- b) The root object and all partitions.

If after a hard reset the device server has determined that processing of an OBJECT STRUCTURE CHECK command for the root object and all partitions is necessary to ensure proper OBSD storage integrity, then it shall:

- a) Terminate all received commands except INQUIRY, REPORT LUNS, and REQUEST SENSE with CHECK CONDITION status, with the sense key set to NOT READY, the additional sense code set to LOGICAL UNIT NOT READY, STRUCTURE CHECK REQUIRED, and the INFORMATION field set to zero; and
- b) Complete received REQUEST SENSE commands with GOOD status, with the sense key set to NOT READY, the additional sense code set to LOGICAL UNIT NOT READY, STRUCTURE CHECK REQUIRED, and the INFORMATION field set to zero.

The need to process an OBJECT STRUCTURE CHECK command for the root object and all partitions shall not affect the processing of the INQUIRY command and REPORT LUNS command.

If after a hard reset the device server has determined that processing of an OBJECT STRUCTURE CHECK command for a partition is necessary to ensure proper OBSD storage integrity, then it shall terminate all received commands addressed to that partition with CHECK CONDITION status, with the sense key set to NOT READY, the additional sense code set to LOGICAL UNIT NOT READY, STRUCTURE CHECK REQUIRED, and the INFORMATION field set to the Partition_ID of a partition for which the processing of an OBJECT STRUCTURE CHECK command is needed.

The status and sense data described in this subclause is returned for all received commands until a suitable OBJECT STRUCTURE COMMAND command has begun processing.

4.12 Security

4.12.1 Basic security model

The OSD security model is a credential-based access control system composed of the following components:

- a) An OBSD (see 3.1.27);
- b) A policy/storage manager (see 4.11);
- c) A security manager; and
- d) Application clients.

The principal function of the security manager is preparing credentials (see 3.1.11) in response to application client requests. A credential is a data structure containing one or more capabilities prepared by the policy/storage manager (see 4.11) and protected by one or two integrity check values (see 3.1.19), having the following properties:

- a) The capability or capabilities (see 3.1.4) in the credential grant defined access to an OSD logical unit for specific command functions (see 3.1.10); and
- b) The integrity check values in the credential protect the capabilities and commands that include the capabilities from various attacks described in 4.12.4.

Figure 4 shows the flow of transactions between the components of the OSD security model.

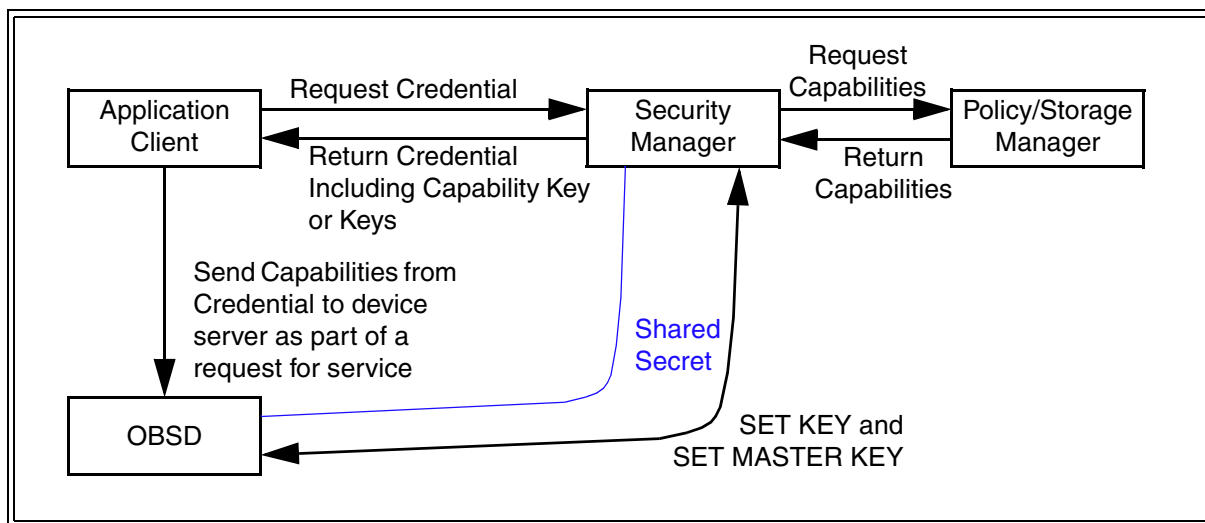


Figure 4 — OSD security model transactions

The security manager generates credentials, including capabilities prepared by the policy/storage manager, for authorized application clients at the request of an application client. The security manager returns one or two capability keys with each credential. The credential gives the application client access to specific OSD components. The capability keys allow the application client and device server to authenticate the commands and data they exchange with an integrity check value (see 4.12.8).

The protocol between the application client and the security manager is not defined by this standard. However, the structure of the credential returned from the security manager to the application client is.

If any security method except NOSEC is used, the device server validates each command received from an application client to confirm that:

- a) The credential has not been tampered with (i.e., that the credential was generated by the security manager and includes an integrity check value using a secret key known only to the security manager and OSD device server); and
- b) The credential was rightfully obtained by the application client from the security manager or through delegation by another application client (i.e., that the application client knows the capability keys that are associated with the credential and has used the capability keys to provide a proper integrity check value or values for the command); and
- c) The requested command function is permitted by the capability or capabilities in the credential as described in 4.11.2.

The capability keys allow the OSD device server to validate that an application client rightfully obtained a credential and that the capability or capabilities have not been tampered with. An application client that has just the capability (e.g., obtained by monitoring CDBs sent to the OSD device server) but not the associated capability key or keys unable to generate commands with valid integrity check value, meaning that application client is denied access to the OSD logical unit. This protocol allows delegation of a credential if an application client delegates both the credential and the capability keys.

The application client requests credentials and capability keys from the security manager for the command functions it needs to perform and sends those capabilities in those credentials to the OSD device server as part of commands that include an integrity check value using the capability keys. While the application client is not trusted to follow this protocol, an application client that does not follow the protocol is unlikely to receive service from the OSD device server.

The security manager may authenticate the application client, but the OSD device server does not authenticate the application client. It is sufficient for the OSD device server to verify the capabilities and integrity check values sent by the application client.

4.12.2 Trust assumptions

This subclause describes how each component of the OSD security model trusts the other components.

The OBSD is a trusted component, meaning that once an application client authenticates that it is communicating with a specific OSD logical unit using methods outside the scope of this standard, it trusts the OBSD to:

- a) Provide integrity for stored data;
- b) Perform the security protocol and functions defined for it by this standard; and
- c) Not be controlled in a way that operates to the detriment of the application client's interests.

The security manager is a trusted component. After the security manager is authenticated by the application client and by the OBSD using methods outside the scope of this standard, the security manager is trusted to:

- a) Safely store long-lived keys;
- b) In cooperation with the policy/storage manager (see 4.11), apply access controls correctly according to requirements that are outside the scope of this standard;
- c) Perform the security functions defined for it by this standard; and
- d) Not be controlled in a way that operates to the detriment of the application client's or OSD logical unit's interests.

The application client is not a trusted component. However, the OSD security model is defined so that the application client receives service from the OSD device server only if it interacts with both the security manager and the OSD device server in ways that assure the propriety of the application client's actions.

The OSD security model components are trusted to protect capability keys from disclosure to unauthorized entities.

The OSD security model components are trusted to maintain some degree of synchronization between their clocks. The OSD security model includes features designed to manage the dependency on the degree of clock synchronization maintained by application clients (see 4.12.7).

Regardless of where the security manager resides (see 4.4), communications between the security manager and other components are trusted based on the requirements shown in table 28.

Table 28 — Security manager communications trust requirements

Component	Security Manager communications trust requirement
OSD device server	Same as for any application client
Application client	Confidential ^a
Policy/storage manager	Message Integrity ^b
^a Confidential communications shall be protected from eavesdropping by physical or cryptographic means. ^b Message integrity assures that the message received is the one that was sent (i.e., no tampering occurred). Messages in which tampering is detected are discarded.	

4.12.3 Preparing credentials

4.12.3.1 Introduction

In response to a request from an application client, the security manager shall prepare and return a credential as follows:

- 1) Forward the access requests from the application client to the policy/storage manager. If the policy/storage manager denies the forwarded requests an error shall be returned to the requesting application client;
- 2) Insert one of the capabilities returned by the policy/storage manager as the solo capability (see 3.1.46) in the credential;
- 3) Set the credential OSD SYSTEM ID field to the value in the OSD system ID attribute in the Root Information attributes page (see 7.1.3.8) of the OSD logical unit to which the credential applies;
- 4) Setup the solo capability as described in 4.12.3.2;
- 5) Unless the SECURITY METHOD field in the solo capability contains NOSEC, compute the solo credential integrity check value as described in 4.12.6.3.1 and place the result in the SOLO CREDENTIAL INTEGRITY CHECK VALUE field;
- 6) If more than one capability was requested from and returned by the policy/storage manager, then:
 - 1) Insert each capability that is not the solo capability as an extension capability;
 - 2) Setup each extension capability as described in 4.12.3.2;
 - 3) Compute the total bytes of extension capabilities and place this value in the EXTENSION CAPABILITIES LENGTH field; and
 - 4) Unless the SECURITY METHOD field in the solo capability contains NOSEC, compute the extended credential integrity check value as described in 4.12.6.3.2 and place the result in the EXTENDED CREDENTIAL INTEGRITY CHECK VALUE field;
- 7) If only one capability was requested from and returned by the policy/storage manager, set the EXTENSION CAPABILITIES LENGTH field to zero;
- 8) If the SECURITY METHOD field in the solo capability contains NOSEC, set the SOLO CREDENTIAL INTEGRITY CHECK VALUE field and the EXTENDED CREDENTIAL INTEGRITY CHECK VALUE field to zero; and
- 9) Return the credential thus constructed to the application client.

4.12.3.2 Capability setup steps for credential preparation

For each capability (see 4.11.2) in a credential, the security manager shall setup the capability as follows:

- 1) Set the capability SECURITY METHOD field as follows:
 - A) Select a security method other than the partition default:
 - a) If the application client requested use of a specific security method, and use of the requested security method is allowed by both the addressed partition and the maintained security policy information, set the capability SECURITY METHOD field to the requested value; or
 - b) If the maintained security policy information requires use of a specific security method for the requesting application client, set the capability SECURITY METHOD field to that value;
 - or
 - B) Use the partition default:
 - a) If the application client requested a credential to be used in a SET KEY command (see 6.37) or a SET MASTER KEY command (see 6.38), set the capability SECURITY METHOD field to the value in the default security method attribute in the Root Policy/Security attributes page (see 7.1.3.22);
 - b) If the capability OBJECT TYPE field contains ROOT, set the capability SECURITY METHOD field to the value in the default security method attribute in the Root Policy/Security attributes page;
 - c) If the capability OBJECT TYPE field contains PARTITION, set the capability SECURITY METHOD field to the value in the default security method attribute in the Partition Policy/Security attributes page (see 7.1.3.23) for partition zero (see 4.6.4); or
 - d) Otherwise, set the capability SECURITY METHOD field to the value in the default security method attribute in the Partition Policy/Security attributes page for the partition whose Partition ID is contained in the capability ALLOWED PARTITION_ID field;
- 2) If the SECURITY METHOD field contains NOSEC, then exit the capability setup steps for the preparation of this credential described in this subclause, otherwise;
- 3) Set the capability KEY VERSION field to the number of the working key secret key used to compute the applicable integrity check value in the credential. If a secret key other than a working key is used by the integrity check value computation (e.g., for a SET KEY command (see 6.37) or a SET MASTER KEY command (see 6.38)), then set the capability KEY VERSION field to zero;
- 4) Set the capability INTEGRITY CHECK VALUE ALGORITHM field to the low order four bits of the attribute number of the attribute in the Root Policy/Security attributes page (see 7.1.3.22) that indicates the algorithm used to compute all integrity check values related to this capability (e.g., if attribute number 8000 0003h identifies the integrity check value algorithm used in this capability, then the INTEGRITY CHECK VALUE ALGORITHM field shall contain three); and
- 5) As specified by the maintained security policy information, modify other capability fields, including but not limited to the following:
 - A) Setting the CAPABILITY EXPIRATION TIME field to a value that is consistent with the policy;
 - B) Ensuring that the capability AUDIT field and CAPABILITY DISCRIMINATOR field contain non-zero values;
 - C) Setting the capability OBJECT CREATED TIME field to a non-zero value that is consistent with 4.11.2.2.1 usage; and
 - D) Ensuring that the POL/SEC bit in the PERMISSIONS BIT MASK field is set to zero, if appropriate.

Successful use of the capability expiration time (see item A) in step 5)) requires some degree of synchronization between the clocks of the device server and security manager. The protocol for synchronizing the clocks is outside the scope of this standard, however, the protocol should be implemented in a secure manner (e.g., it should not be possible for an adversary to set the clock in the device server backwards to enable the reuse of expired credentials).

4.12.4 Security methods

4.12.4.1 Introduction

This standard defines several security methods (see table 29).

Table 29 — OSD security methods

Security Method	Description	Security Method coded value ^a	Reference
NOSEC	No security	0h	4.12.4.2
CAPKEY	Integrity of capabilities	1h	4.12.4.3
CMDRSP	Integrity of CDB, status, and sense data	2h	4.12.4.4
ALLDATA	Integrity of all data in transit	3h	4.12.4.5
^a Security method coded values are used in the capability SECURITY METHOD field and least significant four bits of default security method attributes (e.g., the default security method attribute in the Partition Policy/Security attributes page (see 7.1.3.23)). Security method values 4h to Fh are reserved.			

The security method used by one partition may be different from the security method used by another partition.

A command prepared for a security mode other than the one specified in the CDB SECURITY METHOD field may complete without errors (e.g., a command prepared for the ALLDATA security method may complete without errors reported by the device server if the CMDRSP security method is in use because the preparations for the ALLDATA security method include the preparations that are necessary for the CMDRSP security method).

The OSD security methods are designed to address zero or more specific security threats (see table 30).

Table 30 — Security methods and threats thwarted

Threat	Threat thwarted by security method				
	NOSEC	CAPKEY		CMDRSP	ALLDATA
		Over secure channel ^a			
		No	Yes		
Forgery of credential	No	Yes	Yes	Yes	Yes
Alteration of capabilities	No	Yes	Yes	Yes	Yes
Use of credential by unauthorized application client	No	Yes ^b	Yes ^c	Yes	Yes
Replay of command or status	No	No	Yes ^c	Yes	Yes
Alteration of command or status	No	No	Yes ^c	Yes	Yes
Replay of data	No	No	Yes ^c	No	Yes
Alteration of data	No	No	Yes ^c	No	Yes
Inspection of command, status or data	No	No	Yes/No ^d	No	No
^a This model assumes that one secure channel supports no more than one I_T nexus and that I_T nexus is not shared by multiple application clients. If a SCSI initiator device allows multiple application clients to share an I_T nexus, then the SCSI initiator device implementation and/or application clients shall provide security guarantees equivalent to those provided by a secure channel.					
^b If more than one application client has access to an I_T nexus, then credentials are not protected from use by unauthorized application clients.					
^c A secure channel provides the following security guarantees: a) Cryptographic integrity: Any message received is the one was sent (i.e., no tampering occurred). Messages in which tampering is detected are discarded; b) Data origin authentication: The message received originated from the authenticated originator within the limits of the secure channel authentication mechanism; and c) Replay protection: The same message is not delivered multiple times and that there is a limited number of out-of-order messages.					
^d Optionally, a secure channel may provide a Data Confidentiality guarantee that if a message is read, it cannot be understood other than by the unauthorized parties.					

4.12.4.2 The NOSEC security method

In the NOSEC security method, no OSD security features or algorithms are used by the device server. If the root object and all partitions in the OSD logical unit use the NOSEC security method, then:

- a) Specific SPC-3 commands (e.g., LOG SENSE) may be sent (see table 80 in 6.1) without encapsulating them in the PERFORM SCSI COMMAND command (see 6.23); and
- b) Persistent reservations (see 4.17) are allowed for the logical unit.

4.12.4.3 The CAPKEY security method

The CAPKEY security method validates the integrity of the capability information in each CDB and in the CDB continuation segment (see 5.3), if any.

The application client computes the CDB REQUEST INTEGRITY CHECK VALUE field (see 5.2.11) contents using:

- a) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the solo capability INTEGRITY CHECK VALUE ALGORITHM field (see 4.12.3);
- b) The security token returned in the Security Token VPD page (see 7.5.3); and
- c) The capability key (see 4.12.5.2) for the solo capability.

If the CDB CONTINUATION LENGTH field (see 5.2.1) is not set to zero, the application client computes the CONTINUATION INTEGRITY CHECK VALUE field contents in the CDB continuation segment (see 5.3) using:

- a) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the solo capability INTEGRITY CHECK VALUE ALGORITHM field (see 4.12.3);
- b) The security token returned in the Security Token VPD page (see 7.5.3); and
- c) The applicable capability key (see 4.12.5.2).

The device server validates the credential as described in 4.12.6.1.

The CAPKEY security method is useful when the service delivery subsystem between the OSD device server and application client is secured via methods specified in the applicable SCSI transport protocol, with both the CAPKEY security method and SCSI transport protocol secure channel contributing to securing communications as shown in table 30 (see 4.12.4.1).

4.12.4.4 The CMDRSP security method

The CMDRSP security method validates the integrity of the CDB, status, and sense data for each command.

The application client computes the CDB REQUEST INTEGRITY CHECK VALUE field (see 5.2.11) contents using:

- a) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the solo capability INTEGRITY CHECK VALUE ALGORITHM field (see 4.12.3);
- b) All the bytes in the CDB with the bytes in the REQUEST INTEGRITY CHECK VALUE field set to zero; and
- c) The capability key (see 4.12.5.2) for the solo capability.

If the CDB CONTINUATION LENGTH field (see 5.2.1) is not set to zero, the application client computes the CONTINUATION INTEGRITY CHECK VALUE field contents in the CDB continuation segment (see 5.3) using:

- a) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the solo capability INTEGRITY CHECK VALUE ALGORITHM field (see 4.12.3);
- b) The following array of bytes:
 - 1) All the bytes in the CDB REQUEST INTEGRITY CHECK VALUE field; and
 - 2) All the bytes in the CDB continuation segment (see 5.3) with the bytes in the CONTINUATION INTEGRITY CHECK VALUE field set to zero;
 and
- c) The applicable capability key (see 4.12.5.2).

The device server validates the credential as described in 4.12.6.1.

If the credential validation process validates the request integrity check value without errors and the continuation integrity check value, if any, is validated without errors, then the device server shall:

- 1) Compute an integrity check value for the response data using:
 - A) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the solo capability INTEGRITY CHECK VALUE ALGORITHM field (see 4.12.3);
 - B) The following array of bytes:
 - 1) The request nonce from the CDB (see 5.2.11);
 - 2) The status byte; and
 - 3) If the status is CHECK CONDITION, the sense data with the RESPONSE INTEGRITY CHECK VALUE field in the OSD response integrity check value sense data descriptor (see 4.16.2.2) set to zero; and
 - C) The applicable capability key (see 4.12.5.2) for the reconstructed credential (see 4.12.6.2); and
- 2) Place the computed integrity check value in the following location:
 - A) If the status is not CHECK CONDITION, the computed integrity check value shall be placed in the response integrity check value attribute in the Current Command attributes page (see 7.1.3.31); or
 - B) If the status is CHECK CONDITION, the computed integrity check value shall be placed in the RESPONSE INTEGRITY CHECK VALUE field in the OSD response integrity check value sense data descriptor (see 4.16.2.2) in the sense data.

If the credential validation process fails to validate the integrity check value associated with the command, the device server shall place zero in the RESPONSE INTEGRITY CHECK VALUE field in the OSD response integrity check value sense data descriptor in the sense data.

If the status is not CHECK CONDITION, the application client validates the response integrity check value by recomputing it as described in this subclause and comparing the result to the value of the response integrity check value attribute in the Current Command attributes page.

If the status is CHECK CONDITION, the application client validates the response integrity check value by:

- 1) Saving the response integrity check value found in the RESPONSE INTEGRITY CHECK VALUE field in the OSD response integrity check value sense data descriptor in the sense data;
- 2) Placing zero in the RESPONSE INTEGRITY CHECK VALUE field in the OSD response integrity check value sense data descriptor (see 4.16.2.2);
- 3) Recomputing the response integrity check value as described in this subclause; and
- 4) Comparing the result to the value saved in step 1).

If the application client fails in validating the response integrity check value as described in this subclause, it should take a recovery action not specified by this standard (e.g., one possible action is to request a new credential from the security manager and retry the command). If the error reoccurs, alternate recovery actions should be considered and the presence of malicious entities perpetrating a denial of service attack should be considered.

The CMDRSP security method may be used when the service delivery subsystem between the OSD device server and application client is not secured. The CMDRSP security method protects against corruption of the command command parameter data, status, and sense data while avoiding the overhead that may be required to protect all transferred data. Use of the CMDRSP security method prevents an untrusted application client from forging, modifying or replaying a capability.

4.12.4.5 The ALLDATA security method

The ALLDATA security method validates the integrity of all data in transit between an application client and device server.

The application client computes the CDB REQUEST INTEGRITY CHECK VALUE field (see 5.2.11) contents using the same algorithm specified for the CMDRSP security method (see 4.12.4.4). The device server validates the credential as described in 4.12.6.1. If the CDB CONTINUATION LENGTH field (see 5.2.1) is not set to zero, application client computes the CDB CONTINUATION INTEGRITY CHECK VALUE field (see 5.3) contains using the same algorithm specified for the CMDRSP security method (see 4.12.4.4).

The application client also computes the data-out integrity check value using:

- a) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the solo capability INTEGRITY CHECK VALUE ALGORITHM field (see 4.12.3);
- b) The following array of bytes:
 - 1) All the bytes in the CDB REQUEST INTEGRITY CHECK VALUE field (see 5.2.11); and
 - 2) The used bytes in the following Data-Out Buffer segments (see 4.15.4):
 - 1) Command data;
 - 2) Set attributes; and
 - 3) Get attributes;
 and
- c) The applicable capability key (see 4.12.5.2).

The application client places the data-out integrity information (see table 31) in the Data-Out Buffer starting at the byte specified by the CDB DATA-OUT INTEGRITY CHECK VALUE OFFSET field (see 5.2.11).

Table 31 — Data-out integrity information format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
7	NUMBER OF COMMAND DATA BYTES _____ (LSB)							
8	(MSB) _____							
15	NUMBER OF SET ATTRIBUTES BYTES _____ (LSB)							
16	(MSB) _____							
23	NUMBER OF GET ATTRIBUTES BYTES _____ (LSB)							
24	(MSB) _____							
55	DATA-OUT INTEGRITY CHECK VALUE _____ (LSB)							

The NUMBER OF COMMAND DATA BYTES field specifies the number of bytes from the command data segment that are included in the data-out integrity check value. If the value in the CDB LENGTH field, if any, is larger than the value in the NUMBER OF COMMAND DATA BYTES field, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The NUMBER OF SET ATTRIBUTES BYTES field specifies the number of bytes from the set attributes segment that are included in the data-out integrity check value. If the value in the CDB SET ATTRIBUTE LENGTH field, if any, or the value in the CDB SET ATTRIBUTES LIST LENGTH field, if any, is larger than the value in the NUMBER OF SET ATTRIBUTES BYTES field, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The NUMBER OF GET ATTRIBUTES BYTES field specifies the number of bytes from the get attributes segment that are included in the data-out integrity check value. If the value in the CDB GET ATTRIBUTES LIST LENGTH field, if any, is larger than the value in the NUMBER OF GET ATTRIBUTES BYTES field, the command shall be terminated with CHECK

CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The DATA-OUT INTEGRITY CHECK VALUE field contains the data-out integrity check value computed by the application client.

The device server shall validate the data-out integrity check value by:

- 1) Computing an integrity check value using:
 - A) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the solo capability INTEGRITY CHECK VALUE ALGORITHM field (see 4.12.3);
 - B) The following array of bytes:
 - 1) All the bytes in the CDB REQUEST INTEGRITY CHECK VALUE field (see 5.2.11); and
 - 2) The following bytes from Data-Out Buffer:
 - 1) The number of bytes specified by the NUMBER OF COMMAND DATA BYTES field starting at the Data-Out Buffer byte that follows the CDB continuation segment (see 4.15.4) (i.e., the byte offset specified by the contents of the CDB CONTINUATION LENGTH field (see 5.2.1);
 - 2) The number of bytes specified by the NUMBER OF SET ATTRIBUTES BYTES field starting at the Data-Out Buffer byte offset specified by the CDB SET ATTRIBUTES LIST OFFSET field (see 5.2.6.4); and
 - 3) The number of bytes specified by the NUMBER OF GET ATTRIBUTES BYTES field starting at the Data-Out Buffer byte offset specified by the CDB GET ATTRIBUTES LIST OFFSET field (see 5.2.6.4);
 - and
 - C) The applicable capability key (see 4.12.5.2) for the reconstructed credential (see 4.12.6.2); and
- 2) Comparing the results to contents of the DATA-OUT INTEGRITY CHECK VALUE field.

If the validation fails, the state of the OSD objects and attributes shall not be altered in any detectable way, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID DATA-OUT BUFFER INTEGRITY CHECK VALUE.

The device server shall compute the response integrity check value using the same algorithm specified for the CMDRSP security method (see 4.12.4.4) and the application client validates the response integrity check value using the same algorithm specified for the CMDRSP security method.

The device server shall compute the data-in integrity check value using:

- a) The algorithm indicated by the attribute in the Root Policy/Security attributes page whose attribute number is specified in the solo capability INTEGRITY CHECK VALUE ALGORITHM field;
- b) The following array of bytes:
 - 1) All the bytes in the CDB REQUEST INTEGRITY CHECK VALUE field (see 5.2.11); and
 - 2) The used bytes in the following Data-In Buffer segments (see 4.15.3):
 - 1) Command data or parameter data; and
 - 2) Retrieved attributes;
- and
- c) The applicable capability key (see 4.12.5.2) for the reconstructed credential (see 4.12.6.2).

The device server shall place the data-in integrity information (see table 32) in the Data-In Buffer starting at the byte specified by the CDB DATA-IN INTEGRITY CHECK VALUE OFFSET field (see 5.2.11).

Table 32 — Data-in integrity information format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
7	NUMBER OF COMMAND OR PARAMETER BYTES (LSB)							
8	(MSB) _____							
15	NUMBER OF RETRIEVED ATTRIBUTES BYTES (LSB)							
16	(MSB) _____							
47	DATA-IN INTEGRITY CHECK VALUE (LSB)							

The NUMBER OF COMMAND OR PARAMETER BYTES field specifies the number of bytes from the command data or parameter data segment that are included in the data-in integrity check value.

The NUMBER OF RETRIEVED ATTRIBUTES BYTES field specifies the number of bytes from the retrieved attributes segment that are included in the data-in integrity check value.

The DATA-IN INTEGRITY CHECK VALUE field contains the data-in integrity check value computed by the device server.

After status has been received, the application client validates the data-in integrity check value by:

- 1) Computing an integrity check value using:
 - A) The algorithm indicated by the attribute in the Root Policy/Security attributes page whose attribute number is specified in the solo capability INTEGRITY CHECK VALUE ALGORITHM field;
 - B) The following array of bytes:
 - 1) All the bytes in the CDB REQUEST INTEGRITY CHECK VALUE field (see 5.2.11); and
 - 2) The following bytes from Data-In Buffer:
 - 1) The number of bytes specified by the NUMBER OF COMMAND OR PARAMETER BYTES field starting at the Data-In Buffer byte offset zero; and
 - 2) The number of bytes specified by the NUMBER OF RETRIEVED ATTRIBUTES BYTES field starting at the Data-In Buffer byte offset specified by the CDB RETRIEVED ATTRIBUTES OFFSET field (see 5.2.6);
 - and
 - C) The applicable capability key (see 4.12.5.2);
 - and
- 2) Comparing the results to contents of the DATA-IN INTEGRITY CHECK VALUE field.

If the application client fails in validating the data-in integrity check value, it should take a recovery action not specified by this standard (e.g., one possible action is to request a new credential from the security manager and retry the command). If the error reoccurs, alternate recovery actions should be considered and the presence of malicious entities perpetrating a denial of service attack should be considered.

The ALLDATA security method provides for applying integrity check values to every byte exchanged between the application client and OSD device server. Protection is provided against network attacks similar to those protected against by the security architecture for the internet protocol when confidentiality is not used (see RFC 2401), at the expense of computing and validating numerous integrity check values.

4.12.5 Credentials

4.12.5.1 Credential format

A credential (see table 33) is transferred from the security manager to an application client over a communications mechanism that meets the requirements specified in 4.12.2.

Table 33 — Credential format

Bit Byte	7	6	5	4	3	2	1	0
0	Solo capability (see 4.11.2.2)							
103								
104	OSD SYSTEM ID							
123								
124	(MSB)	SOLO CREDENTIAL INTEGRITY CHECK VALUE						(LSB)
155								
156	(MSB)	EXTENSION CAPABILITIES LENGTH (k-159)						(LSB)
159								
160	Extension capability (see 4.11.2.2)							
263								
	⋮							
k-103	Extension capability (see 4.11.2.2)							
k								
k+1	(MSB)	EXTENDED CREDENTIAL INTEGRITY CHECK VALUE						(LSB)
k+32								

The solo capability is a capability (see 4.11.2.2) to be copied to a CDB (see 5.2.1).

The OSD SYSTEM ID field specifies the value in the OSD system ID attribute in the Root Information attributes page (see 7.1.3.8) of the OSD logical unit to which the credential applies.

The SOLO CREDENTIAL INTEGRITY CHECK VALUE field contains an integrity check value (see 4.12.8) that is computed using the algorithm, inputs, and secret key specified in 4.12.6.3.1.

The EXTENSION CAPABILITIES LENGTH field specifies the number of bytes that follow in zero or more extension capabilities.

Each extension capability is a capability (see 4.11.2.2). All the extension capabilities in a credential are copied to an extension capability CDB continuation descriptor (see 5.4.6).

The EXTENDED CREDENTIAL INTEGRITY CHECK VALUE field contains an integrity check value (see 4.12.8) that is computed using the algorithm, inputs, and secret key specified in 4.12.6.3.2. If the EXTENSION CAPABILITIES LENGTH field is set to zero, then the EXTENDED CREDENTIAL INTEGRITY CHECK VALUE field should be set to zero.

4.12.5.2 Capability keys

All security methods except the NOSEC security method require the computation of one or more integrity check values using one or both of the integrity check values in the credential (see 4.12.5.1) as a capability key secret key (see 3.1.40) as follows:

- a) If the CDB CONTINUATION LENGTH field (see 5.2.1) contains zero (i.e., if there is no CDB continuation segment), then contents of the SOLO CREDENTIAL INTEGRITY CHECK VALUE field are the only capability key used to validate all aspects of the command; or
- b) If the CDB CONTINUATION LENGTH field contains a non-zero value, then:
 - A) If the CDB continuation segment (see 5.3) contains an extension capabilities CDB continuation descriptor (see 5.4.6), then the integrity check values in the credential are used as follows:
 - a) The contents of the SOLO CREDENTIAL INTEGRITY CHECK VALUE field are the capability key that is used to validate the CDB; and
 - b) The contents of the EXTENDED CREDENTIAL INTEGRITY CHECK VALUE field are the capability key that is used to validate all other aspects of the command;
 or
 - B) If the CDB continuation segment does not contain an extension capabilities CDB continuation descriptor, then the contents of the SOLO CREDENTIAL INTEGRITY CHECK VALUE field are the capability key that is used to validate all aspects of the command.

The device server processing of each command relies on only the capability portion or portions of the credential (see 4.12.5.1) that the application client has copied into the CDB and CDB continuation segment (see 5.3). Since the capability or capabilities do not include the integrity check value or values from the credential, the device server needs to compute the capability key or keys for each processed command by:

- 1) Constructing a credential that contains only the solo capability as described in 4.12.6.2.1;
- 2) Computing the solo integrity check value capability key for the constructed credential using the algorithm, inputs, and secret key specified in 4.12.6.3.1; and
- 3) If an extension capabilities CDB continuation descriptor (see 5.4.6) appears in the CDB continuation segment (see 5.3), if any, then:
 - 1) Adding the extension capabilities to the credential constructed in step 1) as described in 4.12.6.2.2; and
 - 2) Computing the extended integrity check value capability key for the constructed credential using the algorithm, inputs, and secret key specified in 4.12.6.3.2.

NOTE 3 The device server may perform the capability key computations steps described in this subclause once for every command processed or repeat them every time a capability key is needed for a validation operation.

4.12.6 OSD device server security algorithms

4.12.6.1 Credential validation

4.12.6.1.1 Introduction

The processes described in this subclause do not apply if the CDB SECURITY METHOD field specifies the NOSEC security method (i.e., if the CDB SECURITY METHOD field contains zero).

If the CDB SECURITY METHOD field specifies the CMDRSP security method or the ALLDATA security method, the device server shall validate the CDB REQUEST NONCE field as described in 4.12.7.2.

The device server validates a credential by verifying the integrity check values in the CDB and CDB continuation segment (see 5.3), if any, that the application client computed using one or both of the capability keys (see 4.12.5.2) in the credential.

The device server shall validate the solo portion of a credential (see 4.12.5) as described in 4.12.6.1.2.

If the CDB CONTINUATION LENGTH field (see 5.2.1) contains a non-zero value, then:

- a) If the CDB continuation segment (see 5.3) contains an extension capabilities CDB continuation descriptor (see 5.4.6), then the device server shall validate the extension portion of a credential as described in 4.12.6.1.3; or
- b) If the CDB continuation segment does not contain an extension capabilities CDB continuation descriptor, then the device server shall revalidate the solo portion of a credential as described in 4.12.6.1.4.

4.12.6.1.2 Validating the solo portion of a credential

The device server shall validate the solo portion of a credential (see 4.12.5) by:

- 1) Constructing a credential that contains only the solo capability as described in 4.12.6.2.1;
- 2) Computing the solo integrity check value capability key for the constructed credential using the algorithm, inputs, and secret key specified in 4.12.6.3.1;
- 3) Computing the request integrity check value using:
 - A) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the capability INTEGRITY CHECK VALUE ALGORITHM field of the solo capability (see 4.12.3);
 - B) Based on the contents of the CDB SECURITY METHOD field, one of the following arrays of bytes:
 - a) For the CAPKEY security method, the security token (see 4.12.4.3); or
 - b) For the CMDRSP security method and the ALLDATA security method, all the bytes in the CDB with the bytes in the REQUEST INTEGRITY CHECK VALUE field set to zero;
 and
 - C) The solo integrity check value capability key computed in step 2) as the secret key;
 and
- 4) Verifying that the request integrity check value matches the contents of the CDB REQUEST INTEGRITY CHECK VALUE field (see 5.2.11). If the contents in the REQUEST INTEGRITY CHECK VALUE field in the CDB do not match the computed solo integrity check value, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the validation of a credential results in CHECK CONDITION status being returned, the state of the OSD objects and attributes shall not be altered in any detectable way.

4.12.6.1.3 Validating the extension portion of a credential

If the CDB continuation segment (see 5.3) contains an extension capabilities CDB continuation descriptor (see 5.4.6), then the device server shall validate the extension portion of a credential (see 4.12.5) by:

- 1) Adding extension capability information to the credential constructed during the validation of the solo portion of the credential (see 4.12.6.1.2) as described in 4.12.6.2.2;
- 2) Computing the extended integrity check value capability key for the constructed and extended credential using the algorithm, inputs, and secret key specified in 4.12.6.3.2;
- 3) Computing the continuation integrity check value using:
 - A) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the capability INTEGRITY CHECK VALUE ALGORITHM field of the solo capability (see 4.12.3);
 - B) Based on the contents of the CDB SECURITY METHOD field, one of the following arrays of bytes:
 - a) For the CAPKEY security method, the security token (see 4.12.4.3); or
 - b) For the CMDRSP security method and the ALLDATA security method, the following array of bytes:
 - 1) All the bytes in the CDB REQUEST INTEGRITY CHECK VALUE field; and

- 2) All the bytes in the CDB continuation segment (see 5.3) with the bytes in the CONTINUATION INTEGRITY CHECK VALUE field set to zero;
- and
- C) The extended integrity check value capability key computed in step 2) as the secret key;
- and
- 4) Verifying that the continuation integrity check value matches the contents of the CONTINUATION INTEGRITY CHECK VALUE field in the CDB continuation segment (see 5.3). If the contents in the CONTINUATION INTEGRITY CHECK VALUE field in the CDB do not match the computed extension integrity check value, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

4.12.6.1.4 Validating a CDB continuation segment using the solo portion of a credential

If the CDB CONTINUATION LENGTH field (see 5.2.1) contains a non-zero value but the CDB continuation segment (see 5.3) does not contain an extension capabilities CDB continuation descriptor (see 5.4.6), then the device server shall validate the CDB continuation segment using the credential's solo portion by:

- 1) Locating or reconstructing the credential constructed in 4.12.6.1.2 that contains only the solo capability and the associated solo integrity check value;
- 2) Computing the continuation integrity check value using:
 - A) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the capability INTEGRITY CHECK VALUE ALGORITHM field of the solo capability (see 4.12.3);
 - B) Based on the contents of the CDB SECURITY METHOD field, one of the following arrays of bytes:
 - a) For the CAPKEY security method, the security token (see 4.12.4.3); or
 - b) For the CMDRSP security method and the ALLDATA security method, the following array of bytes:
 - 1) All the bytes in the CDB REQUEST INTEGRITY CHECK VALUE field; and
 - 2) All the bytes in the CDB continuation segment (see 5.3) with the bytes in the CONTINUATION INTEGRITY CHECK VALUE field set to zero;
 - and
 - C) The solo integrity check value capability key computed in step 1) as the secret key;
 - and
 - 3) Verifying that the continuation integrity check value matches the contents of the CONTINUATION INTEGRITY CHECK VALUE field in the CDB continuation segment (see 5.3). If the contents in the CONTINUATION INTEGRITY CHECK VALUE field in the CDB do not match the computed extension integrity check value, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

4.12.6.2 Reconstructing the credential

4.12.6.2.1 Reconstructing the solo portion of a credential

The device server reconstructs the solo portion of a credential from a CDB capability by:

- a) Copying the value in the OSD system ID attribute in the Root Information attributes page (see 7.1.3.8) to the OSD SYSTEM ID field of the reconstructed credential;
- b) Copying the capability from the CDB to the solo capability portion of the reconstructed credential; and
- c) Setting the EXTENSION CAPABILITIES LENGTH field to zero.

The CREDENTIAL INTEGRITY CHECK VALUE field and EXTENDED CREDENTIAL INTEGRITY CHECK VALUE field are not used in a reconstructed credential and are set to zero.

4.12.6.2 Reconstructing the extension portion of a credential

Using the contents of the reconstructed solo portion of a reconstructed credential (see 4.12.6.2.1), the device server reconstructs the extension portion of a credential from the contents of an extension capabilities CDB continuation descriptor (see 5.4.6) by:

- a) Not modifying the contents of the solo capability portion, OSD SYSTEM ID field, and SOLO CREDENTIAL INTEGRITY CHECK VALUE field in the input reconstructed credential;
- b) Setting the EXTENSION CAPABILITIES LENGTH field in the reconstructed credential to the value in the CDB CONTINUATION DESCRIPTOR LENGTH field of the extension capabilities CDB continuation descriptor minus four; and
- c) Copying all of the bytes in all of the extension capabilities in the extension capabilities CDB continuation descriptor to the extension capabilities portion of the reconstructed credential.

The SOLO CREDENTIAL INTEGRITY CHECK VALUE field and EXTENDED CREDENTIAL INTEGRITY CHECK VALUE field are not used in a reconstructed credential and are set to zero.

4.12.6.3 Computing the integrity check values for a credential

4.12.6.3.1 Computing the solo integrity check value for a credential

The solo credential integrity check value shall be computed using:

- a) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the INTEGRITY CHECK VALUE ALGORITHM field in the solo capability in the credential (see 4.12.3);
- b) All the bytes in:
 - 1) The solo capability portion of the credential (see 4.12.5.1); and
 - 2) The OSD SYSTEM ID field of the credential;
 and
- c) The secret key selected as follows:
 - A) If the OBJECT TYPE field in the credential's (see 4.12.5.1) solo capability (see 4.11.2.2) contains COLLECTION or USER, the secret key is the authentication working key:
 - a) Identified by the KEY VERSION field in the credential's (see 4.12.5.1) solo capability; and
 - b) Associated with the partition identified by the ALLOWED PARTITION_ID field in the credential's (see 4.12.5.1) solo capability;
 - B) If the OBJECT TYPE field in the capability contains ROOT or PARTITION and the command is not SET KEY and not SET MASTER KEY, the secret key is the authentication working key for partition zero identified by the KEY VERSION field in the credential's (see 4.12.5.1) solo capability;
 - C) If the command is SET KEY (see 6.37), the secret key that is selected as follows:
 - a) If the KEY TO SET field in the CDB contains 01b (i.e., update root key), the authentication master key;
 - b) If the KEY TO SET field in the CDB contains 10b (i.e., update partition key), the authentication root key; or
 - c) If the KEY TO SET field in the CDB contains 11b (i.e., update working key), the authentication partition key for the partition identified by the PARTITION_ID field in the CDB;
 or
 - D) For the SET MASTER KEY command:
 - a) For the SEED EXCHANGE step (see 6.38.2), the authentication master key; or
 - b) For the CHANGE MASTER KEY step (see 6.38.3), the next authentication master key computed after GOOD status has been returned by the SEED EXCHANGE step (see 6.38.2).

4.12.6.3.2 Computing the extended integrity check value for a credential

The extended credential integrity check value shall be computed as follows:

- 1) An intermediate integrity check value shall be computed using:
 - A) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the INTEGRITY CHECK VALUE ALGORITHM field in the solo capability in the credential (see 4.12.3);
 - B) All the bytes in the credential (see 4.12.5.1) except the EXTENDED CREDENTIAL INTEGRITY CHECK VALUE field; and
 - C) The secret key selected as follows:
 - a) If the OBJECT TYPE field in the credential's (see 4.12.5.1) solo capability (see 4.11.2.2) contains COLLECTION or USER, the secret key is the authentication working key:
 - A) Identified by the KEY VERSION field in the credential's (see 4.12.5.1) solo capability; and
 - B) Associated with the partition identified by the ALLOWED PARTITION_ID field in the credential's (see 4.12.5.1) solo capability;
 - or
 - b) If the OBJECT TYPE field in the credential's (see 4.12.5.1) solo capability contains ROOT or PARTITION, the secret key is the authentication working key for partition zero identified by the KEY VERSION field in the credential's (see 4.12.5.1) solo capability;
- and
- 2) Each extension capability in the credential (see 4.12.5.1) shall be processed in the order in which it appears in the credential, and a new intermediate integrity check value shall be computed based on the extension capability being processed using:
 - A) The algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the INTEGRITY CHECK VALUE ALGORITHM field in the solo capability in the credential (see 4.12.3);
 - B) All the bytes in the in the previously computed intermediate integrity check value; and
 - C) The secret key selected as follows:
 - a) If the OBJECT TYPE field in the credential's (see 4.12.5.1) extension capability (see 4.11.2.2) being processed contains COLLECTION or USER, the secret key is the authentication working key:
 - A) Identified by the KEY VERSION field in the credential's (see 4.12.5.1) extension capability being processed; and
 - B) Associated with the partition identified by the ALLOWED PARTITION_ID field in the credential's (see 4.12.5.1) extension capability being processed;
 - or
 - b) If the OBJECT TYPE field in the credential's extension capability being processed contains ROOT or PARTITION, the secret key is the authentication working key for partition zero identified by the KEY VERSION field in the credential's (see 4.12.5.1) extension capability being processed.

The extended credential integrity check value is the last intermediate integrity check value computed.

4.12.6.4 Invalidating credentials

The security manager may invalidate the credentials for one OSD object by requesting that the policy/storage manager change the policy access tag attribute in the policy/security attributes page associated with that OSD object or objects (see 4.11.3.2) to a value other than the policy access tag value that is contained in the credential's capability or capabilities.

The security manager may invalidate credentials for an entire partition by using the SET KEY command (see 6.37) to update the working key version used to compute the integrity check values in those credentials.

4.12.7 Request nonces

4.12.7.1 Request nonce format

For some security methods (see 4.12.4), an application client generated request nonce (see table 34) is included in the input data for each integrity check value computation (see 4.12.8) to thwart attempts to capture OSD commands (e.g., FORMAT OSD) and replay them.

Table 34 — Request nonce format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
5	TIMESTAMP							(LSB) _____
6	(MSB) _____							
11	RANDOM NUMBER							(LSB) _____

The **TIMESTAMP** field contains the number of milliseconds that have elapsed since midnight, 1 January 1970 UT (see 3.1.52). Timestamp values should be coordinated with the contents of the clock attribute in the Root Information attributes page (see 7.1.3.8) using techniques that are outside the scope of this standard.

The **RANDOM NUMBER** field contains a random number generated from a good source of entropy (e.g., as described in RFC 1750).

If the security method being used does not require generation of request nonce values, the nonce **TIMESTAMP** field should contain zero.

4.12.7.2 Device server validation of request nonces

If the command is being processed using the **CMDRSP** security method or the **ALLDATA** security method (see 4.12.4) and a request nonce with zero in the **TIMESTAMP** field is received, the command shall be terminated with **CHECK CONDITION** status, with the sense key set to **ILLEGAL REQUEST** and the additional sense code set to **INVALID FIELD IN CDB**.

If the inputs to an integrity check value computation include a non-zero request nonce that is listed (see 4.12.7.3) as having been used in any previous integrity check value computation, the command shall be terminated with **CHECK CONDITION** status, with the sense key set to **ILLEGAL REQUEST** and the additional sense code set to **NONCE NOT UNIQUE**. The command shall be terminated regardless of the success or failure of the previous command in which the duplicate request nonce appeared (e.g., the request nonce appearing in a **WRITE** command that ultimately fails due to insufficient quota or the request nonce appearing in a **CREATE** command that ultimately fails because the computed integrity check value for the credential is wrong shall not be accepted a second time).

If the request nonce timestamp is less than the contents of the clock attribute in the Root Information attributes page (see 7.1.3.8) minus the value in the oldest valid nonce attribute in the Partition Policy/Security attributes page (see 7.1.3.23), then the command shall be terminated with **CHECK CONDITION** status, with the sense key set to **ILLEGAL REQUEST**, and with the additional sense code set to **NONCE TIMESTAMP OUT OF RANGE**. If a command is terminated in this way, the current contents of the clock attribute in the Root Information attributes page shall be returned left-aligned and zero-padded (see 3.8.2) in the **COMMAND-SPECIFIC INFORMATION** field of the command-specific information sense data descriptor.

If the request nonces timestamp is greater than the contents of the clock attribute in the Root Information attributes page plus the value in the newest valid nonce attribute in the Partition Policy/Security attributes page, then the command be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to NONCE TIMESTAMP OUT OF RANGE. If a command is terminated in this way, the current contents of the clock attribute in the Root Information attributes page shall be returned left-aligned and zero-padded (see 3.8.2) in the COMMAND-SPECIFIC INFORMATION field of the command-specific information sense data descriptor.

Successful use of the request nonces requires some degree of synchronization between the clocks of the device server and security manager. The protocol for synchronizing the clocks is outside the scope of this standard, however, the protocol should be implemented in a secure manner (i.e., it should not be possible for an adversary to set the clock in the device server backwards to enable the replay of expired request nonces).

4.12.7.3 Lists of previously used request nonces

4.12.7.3.1 Introduction

The device server shall maintain a list of all request nonces used in integrity check value computations. Failure of the integrity check value computation shall not result in exclusion from the list.

A request nonce shall appear in the list from the time it is received until:

- a) The time in the request nonce timestamp field is less than the value in the clock attribute in the Root Information attributes page (see 7.1.3.8) minus the value in the oldest valid nonce limit attribute in the Root Policy/Security attributes page (see 7.1.3.22);
- b) The working key used to compute the integrity check value in which the request nonce was used is invalidated by a SET KEY command (see 6.37);
- c) Optionally, the capability audit field is frozen (see 4.12.7.3.2); or
- d) Optionally, the working key is frozen (see 4.12.7.3.3).

For the SET KEY command and the SET MASTER KEY command (see 6.38), the request nonce shall appear in the list from the time it is received until the time in the request nonce timestamp field is less than the value in the clock attribute in the Root Information attributes page minus the value in the oldest valid nonce limit attribute in the Root Policy/Security attributes page (i.e., only item a) applies to these commands).

The request nonce list depth attribute in the Root Policy/Security attributes page shall indicate the minimum number of request nonce list entries available to one application client.

4.12.7.3.2 Freezing capability audit fields

The device server may refuse to accept any additional commands containing a specific combination of capability AUDIT field and capability KEY VERSION field values (see 4.11.2.2). If the device server takes this action, it should terminate the selected command and all future commands containing the selected combination of capability AUDIT field and capability KEY VERSION field values with CHECK CONDITION status, a sense key set to ILLEGAL REQUEST, and an additional sense code set to SECURITY AUDIT VALUE FROZEN.

The device server may repeat the process described in this subclause as often as necessary to reduce the amount of resources required to implement the nonce listing requirements (see 4.12.7.3.1).

4.12.7.3.3 Freezing working keys

The device server may refuse to accept any additional commands with a capability KEY VERSION field (see 4.11.2.2) specifying a certain working key version value. If the device server takes this action, it:

- a) Should terminate the selected command and all future commands having the selected capability key version value with CHECK CONDITION status, a sense key set to ILLEGAL REQUEST, and an additional sense code set to SECURITY WORKING KEY FROZEN; and
- b) Shall set to one the bit in the frozen working key bit mask attribute in the Partition Policy/Security attributes page (see 7.1.3.23) that corresponds to the working key version thus selected.

The device server may repeat the process described in this subclause as often as necessary to reduce the amount of resources required to implement the nonce listing requirements (see 4.12.7.3.1).

4.12.8 Integrity check values

An integrity check value is a value produced by a cryptographic function (e.g., HMAC-SHA1) based on a secret key (see 4.12.9) that is able to be computed and verified by the entities knowing the secret key. Integrity check values are used to verify that:

- a) A collection of data fields contain correct values; and
- b) The values in those data fields were prepared by the entity that created the integrity check value.

Some integrity check value algorithms return values that contain fewer bytes than are available in the fields that this standard defines to contain integrity check values. If this occurs, then:

- a) The integrity check value returned by the specified algorithm shall be placed in the field with the most significant byte of the integrity check value being placed in the most significant byte of the field and the remaining integrity check value bytes being placed in consecutive bytes of the field;
- b) Zeros shall be placed in the unused bytes of the field up to and including the least significant byte of the field.

4.12.9 Secret keys

4.12.9.1 Introduction

The hierarchy of secret keys and the mechanisms for updating them are described in:

- a) This subclause;
- b) The definition of the SET MASTER KEY command (see 6.38); and
- c) The definition of the defining the SET KEY command (see 6.37).

In the OSD security model, the security of transactions depends on a hierarchy of secret keys as shown in table 35, with the highest key in the hierarchy (i.e., the master key) shown at the top of the table and the lowest keys in the hierarchy (i.e., the capability keys) shown at the bottom of the table.

Table 35 — OSD secret key hierarchy

Key Name	Key Shared Using	Key Used To	Key Update Frequency
Keys shared between the security manager and the OSD device server			
Master	SET MASTER KEY command	Update Root key	Change of logical unit owner
Root	SET KEY command	Update Partition key	When Partition key may have been compromised (i.e., very infrequently)
Partition ^a	SET KEY command	Update Working keys	When Working key updates may have been compromised (i.e., infrequently)
Working ^b	SET KEY command	Create Capability keys	When normal key use affords too much chance that the working key might be reverse engineered (i.e., regularly)
Keys shared between the security manager and the application client ^c			
Capability ^d	Credentials and mechanisms not specified in this standard	Secure commands, responses, and data	New with each new Credential
^a For the purposes of the secret key hierarchy, the root object is treated the same as any other partition OSD object using partition zero. ^b For each partition, up to sixteen working keys may be active at any time, uniquely identified by the capability KEY VERSION field (see 4.11.2.2). ^c The device server is capable of computing the capability key (see 4.12.6.3) using the reconstructed credential (see 4.12.6.2). ^d As a dual purpose number, the capability key is different from other keys in the hierarchy. A capability key is one of the two integrity check values in a credential (see 4.12.5). Even though the security manager computes it, the computation is based on values beyond the security manager's control (e.g., the user object to which the credential allows access). While changing the working key used to construct integrity check values in a credential invalidates the capability keys, one or more of the capabilities in the credential may expire before that, making those capability keys invalid.			

Each master, root, and partition key represents two secret key values as follows:

- a) An authentication key that is used to compute the credential integrity check values; and
- b) A generation key that is used by future SET KEY commands and SET MASTER KEY commands to compute the updated generation key and new authentication key values.

When an OBSD is manufactured, both the master authentication key and master generation key values shall be provided for each logical unit. The two values may be identical. The initial master keys should be generated as specified by FIPS 198 and the length of initial master keys should comply with FIPS 198.

The secret keys shared between the security manager and OSD device server are very secret information. They should be protected from being discovered by an adversary. They should be stored in a tamper resistant non-volatile manner and may be protected by a tamper resistant software shield. The master key shall be stored in a tamper resistant manner.

The seeds that have been used to create all secret keys other than the master key may be saved in non-volatile memory for later use in recomputing the secret key values. The OSD logical unit should not store the commands sent to set the master key in a manner that has the potential for being externally accessible.

4.12.9.2 Computing updated generation keys and new authentication keys

The SET KEY command (see 6.37) and SET MASTER KEY command (see 6.38) shall perform the steps described in this subclause to compute new generation and authentication keys.

The inputs to the process are:

- a) The input key value is one of the following:
 - A) For a SET KEY command, the generation key from the next higher level in the key hierarchy shall be used (e.g., the root key generation key is used to create the first partition keys for a newly created partition), as selected by the KEY TO SET field in the CDB of that command; or
 - B) For a SET MASTER KEY command, the previous master key generation key shall be used;
- b) The seed value is one of the following:
 - A) For a SET KEY command, the contents of the SEED field of the CDB for the command; or
 - B) For a SET MASTER KEY command key, the value computed after GOOD status has been returned in the SEED EXCHANGE step (see 6.38.2) and updated by CHANGE MASTER KEY step (see 6.38.3);and
- c) The integrity check value algorithm indicated by the attribute in the Root Policy/Security attributes page (see 7.1.3.22) whose attribute number is specified in the capability INTEGRITY CHECK VALUE ALGORITHM field (see 4.12.3) in the capability in the CDB for the command.

The updated generation key shall be computed by performing the specified integrity check algorithm with the following inputs:

- a) Input key value; and
- b) Seed value.

The new authentication key shall be computed by performing the specified integrity check algorithm with the following inputs:

- a) Input key value; and
- b) Seed value with the least significant bit changed as follows:
 - A) If the seed value least significant bit is zero, then it is changed to one; or
 - B) If the seed value least significant bit is one, then it is changed to zero.

4.12.10 OSD security interactions with SPC-3 commands and SAM-3 task management functions

Persistent reservations (see 4.17) are incompatible with an OSD logical unit in which the root object or any partition is using any security method other than NOSEC (see 4.12.4).

Except for the INQUIRY command, the REPORT LUNS command, the REQUEST SENSE command, and the TEST UNIT READY command, all SPC-3 commands are invalid if addressed to an OSD logical unit in which any partition is using any security method other than NOSEC (see table 80 in 6.1). The PERFORM SCSI COMMAND command (see 6.23) allows SPC-3 commands other than persistent reservations commands to be processed under the protection of the current security method.

If the root object or any partition in the OSD logical unit is using any security method other than NOSEC, all SAM-3 task management functions except QUERY TASK shall be ignored and responded to as if they have been successfully processed. The PERFORM TASK MANAGEMENT FUNCTION command (see 6.24) allows SAM-3 task management functions to be processed under the protection of the current security method.

4.13 Object duplication

4.13.1 Overview

The following mechanisms are defined for duplicating the data and attributes contained in one or more user objects and collections in new user objects and collections:

- a) The CREATE SNAPSHOT command (see 4.13.2);
- b) The CREATE CLONE command (see 4.13.2); and
- c) The COPY USER OBJECTS command (see 6.4).

A model for the partition snapshot and clone mechanisms appears in 4.13.2.

The COPY USER OBJECTS command:

- 1) Creates a destination user object; and
- 2) Copies the data and maybe the attributes from a source user object to that destination user object using:
 - A) The object duplication methods described in 4.13.3;
 - B) The object duplication state management methods described in 4.13.4; and
 - C) The object duplication space accounting methods described in 4.13.5.

4.13.2 Snapshot partitions and clone partitions

4.13.2.1 Overview

The following commands create, update, manage, and remove copies of all the user objects, collections, and attributes between two or more partitions:

- a) The CREATE SNAPSHOT command (see 4.13.2.2, 4.13.2.4, and 6.10);
- b) The CREATE CLONE command (see 4.13.2.3 and 6.7);
- c) The REFRESH SNAPSHOT command (see 6.30);
- d) The RESTORE PARTITION FROM SNAPSHOT command (see 6.35);
- e) The DETACH CLONE command (see 4.13.2.5 and 6.12); and
- f) The REMOVE PARTITION command (see 6.34).

In the context of snapshots and clones, the following types of partitions are identified:

- a) Primary (i.e., not a snapshot or a clone);
- b) Snapshot; and
- c) Clone.

Snapshots and clones are partitions that are full copies of a source partition. Other similarities and differences between snapshots and clones are shown in table 36.

Table 36 — Comparison of snapshots and clones

Feature		Partition type		
		Primary	Snapshot	Clone
Allowed to be the source partition for a CREATE SNAPSHOT command (see 6.10)		Yes	No	Yes
Allowed to be the source partition for a CREATE CLONE command (see 6.12)		No	Yes	No
Allowed use in a REFRESH SNAPSHOT command (see 6.30)	Source	Yes	No	Yes
	Destination	No	Yes	No
Allowed use in a RESTORE PARTITION FROM SNAPSHOT command (see 6.35)	Source	No	Yes	No
	Destination	Yes	No	Yes
Allowed to be the partition specified in a REMOVE PARTITION command (see 6.34)		Yes	Yes	Yes
Allowed to be the partition specified in a DETACH CLONE command (see 6.12)		No	No	Yes
Time ordered history (i.e., chain) of partition duplicates maintained		No	Yes	No
The source partition attribute in the Snapshots Information attributes page (see 7.1.3.30) indicates the primary partition from which this partition is descended		n/a	Yes	Yes
Writable in normal (i.e., non error recovery) usage		Yes	No	Yes

4.13.2.2 Snapshot history chains

The device server maintains snapshot forward attribute and the snapshot backward attribute in the Snapshots Information attributes page (see 7.1.3.30) to form a double linked chain of the snapshot partitions descended from a primary or clone partition. The attribute values are:

- a) The Partition_ID (see 4.6.2) of the partition to which the history chain attribute points; or
- b) Zero or an undefined attribute (see 3.1.51) to indicate no history chain linkage exists.

An application client may trace backward or forward in time using these attributes.

Figure 5 shows the history chain for a primary partition with a single snapshot. The conditions show in figure 5 would be present after a CREATE SNAPSHOT command (see 6.10) with partition 1 as the source partition and partition 2 as the destination. Since the contents of snapshot is fixed in time while the primary partition continues to evolve, time might be viewed as flowing in the direction of the arrow.

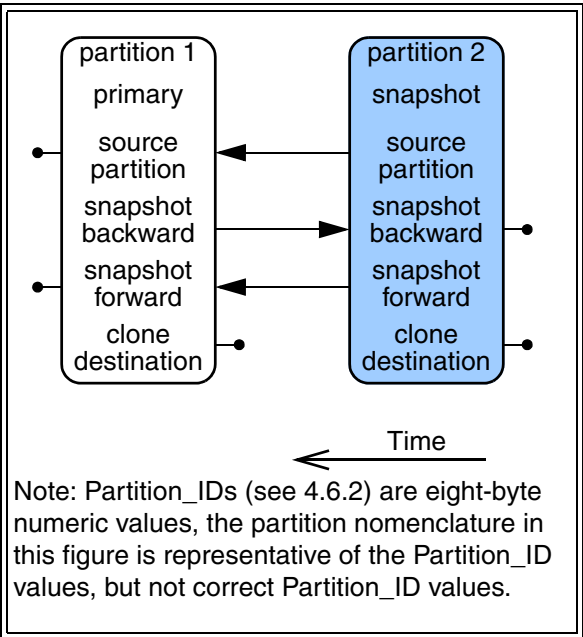


Figure 5 — Snapshot history chain after first CREATE SNAPSHOT command

The attributes in the Snapshots Information attributes page (see 7.1.3.30) that maintain the history chain for figure 5 are summarized in table 37.

Table 37 — Snapshot history chain attributes for one CREATE SNAPSHOT command

Attribute	Partition 1 primary	Partition 2 snapshot
source partition	undefined ^a or zero	partition 1
snapshot backward	partition 2	undefined ^a or zero
snapshot forward	undefined ^a or zero	partition 1
clone destination	undefined ^a or zero	undefined ^a or zero
Note: Partition_IDs (see 4.6.2) are eight-byte numeric values, the partition nomenclature in this table is representative of the Partition_ID values, but not correct Partition_ID values.		
^a See 3.1.51.		

If a second snapshot is taken at a later time, the history chain would become as shown in figure 6.

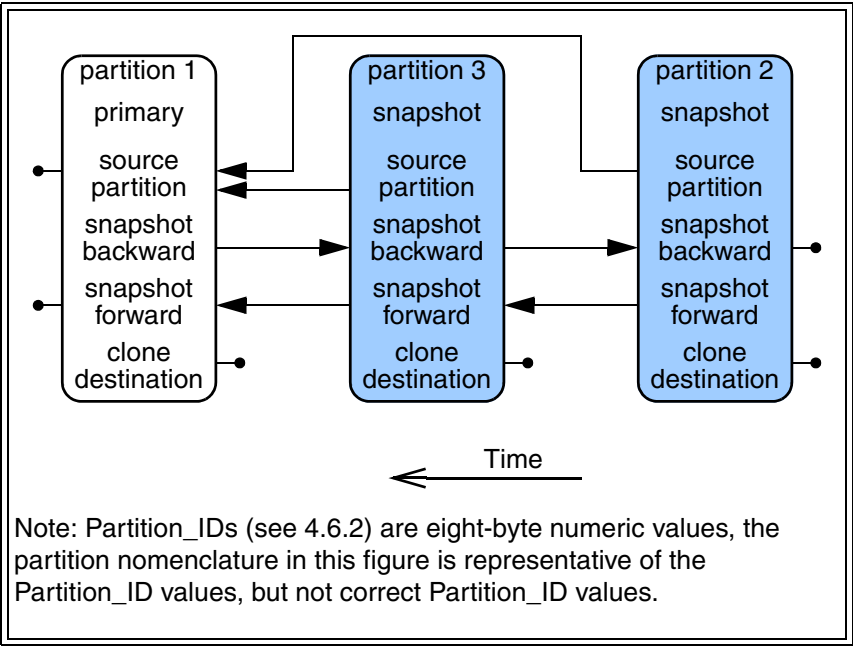


Figure 6 — Snapshot history chain after second CREATE SNAPSHOT command

The attributes in the Snapshots Information attributes page (see 7.1.3.30) that maintain the history chain for figure 6 are summarized in table 38.

Table 38 — Snapshot history chain attributes for two CREATE SNAPSHOT commands

Attribute	Partition 1 primary	Partition 3 snapshot	Partition 2 snapshot
source partition	undefined ^a or zero	partition 1	partition 1
snapshot backward	partition 3	partition 2	undefined ^a or zero
snapshot forward	undefined ^a or zero	partition 1	partition 3
clone destination	undefined ^a or zero	undefined ^a or zero	undefined ^a or zero
Note: Partition_IDs (see 4.6.2) are eight-byte numeric values, the partition nomenclature in this table is representative of the Partition_ID values, but not correct Partition_ID values.			
^a See 3.1.51.			

4.13.2.3 Clone chains

Because clone partitions are writable, they evolve over time in the same way that the partition from which the snapshot was taken evolves. Therefore, no certain temporal relationship is possible with clone partitions.

The clone destination attribute and source partition attribute in the Snapshots Information attributes page (see 7.1.3.30) provide the only linkage between clone partitions and the snapshot partitions from which they are derived.

The presence of multiple clone destination attributes in the Snapshots Information attributes page allows more than one clone partition to be created from a single snapshot partition.

Figure 7 builds on figure 6 (see 4.13.2.2) to show the effects of a CREATE CLONE command (see 6.7) with partition 3 as the source partition and partition 4 as the destination.

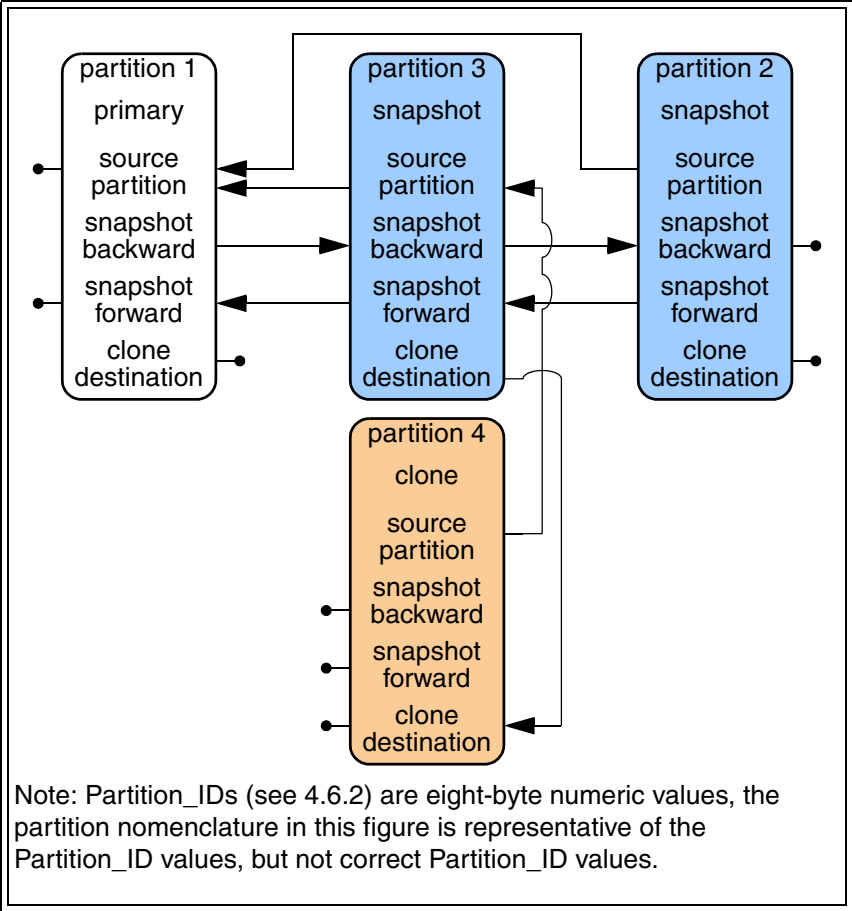


Figure 7 — Snapshot/clone chains after a first CREATE CLONE command

The attributes in the Snapshots Information attributes page (see 7.1.3.30) that maintain the clone chain for figure 7 are summarized in table 39.

Table 39 — Snapshot/clone chain attributes for one CREATE CLONE command

Attribute	Partition 1 primary	Partition 2 snapshot	Partition 3 snapshot	Partition 4 clone
source partition	see table 38 in 4.13.2.2		see table 38 in 4.13.2.2	partition 3
snapshot backward				undefined ^a or zero
snapshot forward				undefined ^a or zero
clone destination				partition 4
Note: Partition_IDs (see 4.6.2) are eight-byte numeric values, the partition nomenclature in this table is representative of the Partition_ID values, but not correct Partition_ID values.				
^a See 3.1.51.				

The effects of a second CREATE CLONE command (see 6.7) with partition 3 as the source are shown in figure 8.

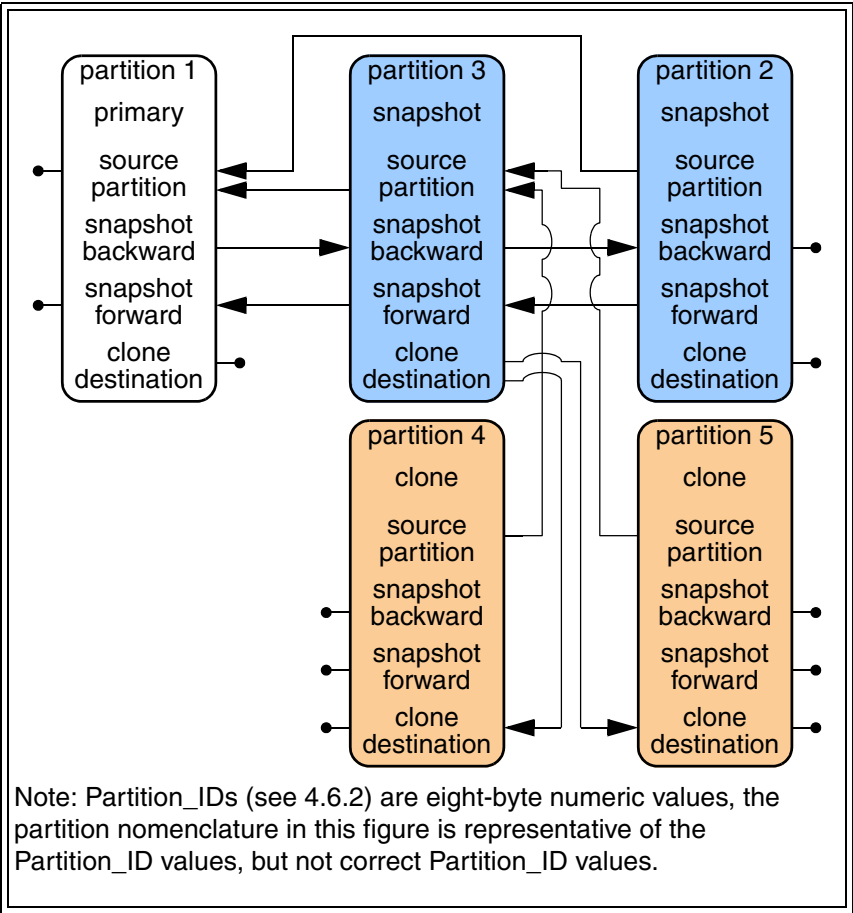


Figure 8 — Snapshot/clone chains after a second CREATE CLONE command

The attributes in the Snapshots Information attributes page (see 7.1.3.30) that maintain the clone chain for figure 8 are summarized in table 40.

Table 40 — Snapshot/clone chain attributes for two CREATE CLONE commands

Attribute	Partition 1 primary	Partition 2 snapshot	Partition 3 snapshot	Partition 4 clone	Partition 5 clone
source partition	see table 38 in 4.13.2.2		see table 38 in 4.13.2.2	partition 3	partition 3
snapshot backward				undefined ^a or zero	
snapshot forward				undefined ^a or zero	
clone destination ^b			partition 4	undefined ^a or zero	
clone destination ^b			partition 5	undefined ^a or zero	
Note: Partition_IDs (see 4.6.2) are eight-byte numeric values, the partition nomenclature in this table is representative of the Partition_ID values, but not correct Partition_ID values.					
^a See 3.1.51.					
^b The Snapshots Information attributes page (see 7.1.3.30) defines several attribute numbers with the name clone destination. Each clone destination attribute points to a different clone partition.					

4.13.2.4 Snapshots of clones

A clone partition may be specified as the source partition for a CREATE SNAPSHOT command (see 6.10). Figure 9 builds on figure 8 (see 4.13.2.3) to show the effects of a CREATE SHAPSHOT command with clone partition 5 as the source partition and partition 6 as the destination.

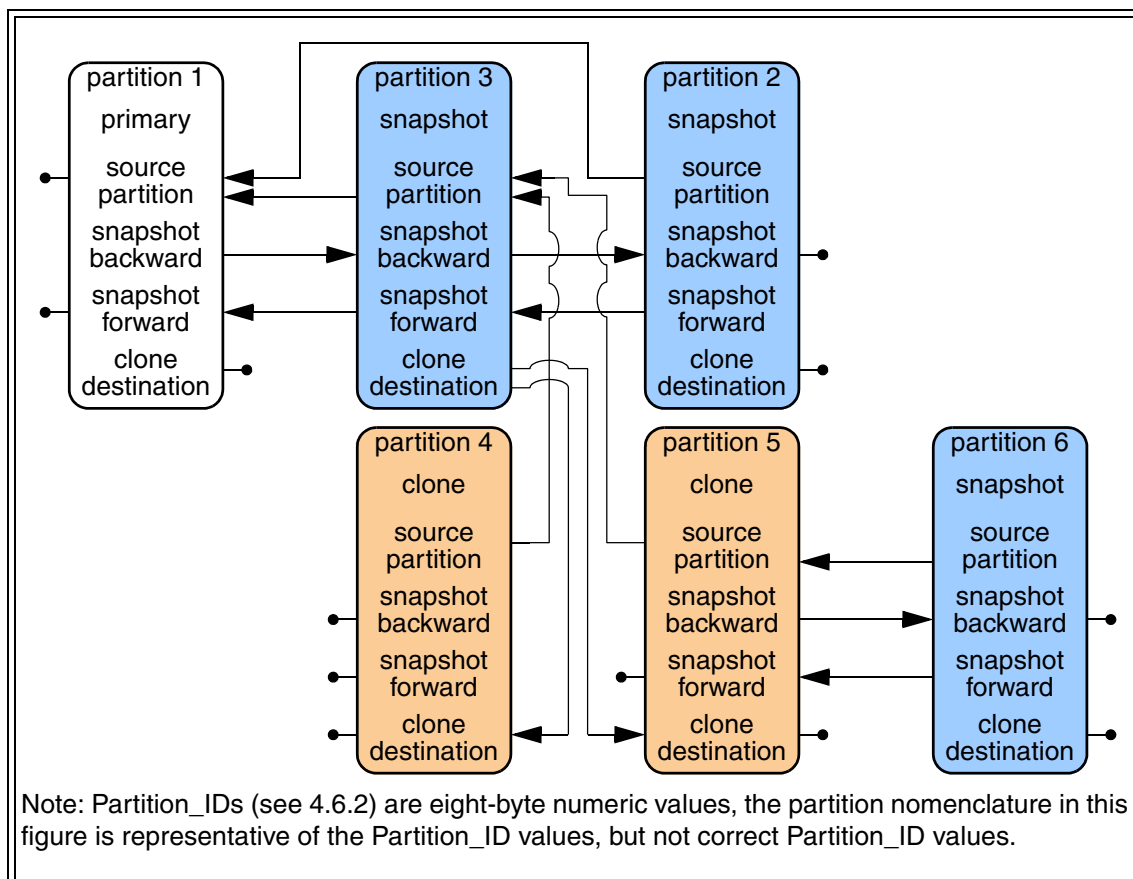


Figure 9 — Snapshot/clone chains after a CREATE SNAPSHOT command on a clone partition

The attributes in the Snapshots Information attributes page (see 7.1.3.30) that maintain the new history chain in figure 9 are summarized in table 41.

Table 41 — Snapshot/clone chain attributes for a CREATE SNAPSHOT command on a clone partition

Attribute	Partition 1 primary	Partitions 2 and 3 snapshots	Partition 4 clone	Partition 5 clone	Partition 6 snapshot
source partition	see table 38 in 4.13.2.2	see table 40 in 4.13.2.3		partition 3	partition 5
snapshot backward				partition 6	undefined ^a or zero
snapshot forward				undefined ^a or zero	partition 5
clone destination				undefined ^a or zero	undefined ^a or zero
Note: Partition_IDs (see 4.6.2) are eight-byte numeric values, the partition nomenclature in this table is representative of the Partition_ID values, but not correct Partition_ID values.					
^a See 3.1.51.					

4.13.2.5 Detaching a clone partition

A clone partition may be detached from its source partition using a DETACH CLONE command (see 6.12). Figure 10 builds on figure 9 (see 4.13.2.4) to show the effects of a DETACH CLONE command for partition 5.

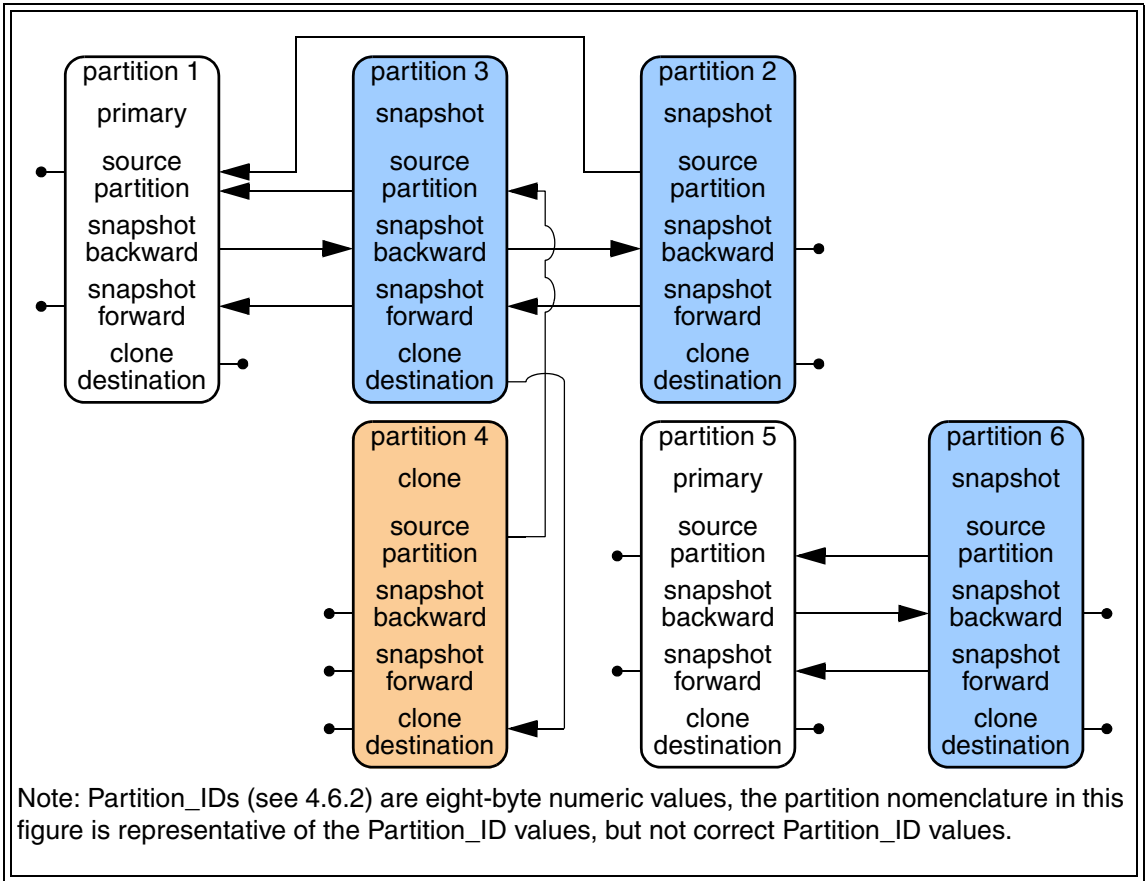


Figure 10 — Snapshot/clone chains after a DETACH CLONE command

The attributes in the Snapshots Information attributes page (see 7.1.3.30) that maintain the two history chains shown in figure 10 are summarized in table 42.

Table 42 — Snapshot/clone chain attributes after a DETACH CLONE command

Attribute	Partition 1 primary	Partitions 2 and 3 snapshots	Partition 4 clone	Partition 5 primary	Partition 6 snapshot
source partition	see table 38 in 4.13.2.2	see table 40 in 4.13.2.3		undefined ^a or zero	partition 5
snapshot backward				partition 6	undefined ^a or zero
snapshot forward				undefined ^a or zero	partition 5
clone destination				undefined ^a or zero	undefined ^a or zero
Note: Partition_IDs (see 4.6.2) are eight-byte numeric values, the partition nomenclature in this table is representative of the Partition_ID values, but not correct Partition_ID values.					
^a See 3.1.51.					

4.13.3 Object duplication methods

Duplicating user object data, collection data, partition data, and attributes may or may not involve making two copies of the same bytes on stable storage (see table 43).

Table 43 — Object duplication methods

Name	Code ^a	Description
DEFAULT	00h	Used to specify one of the other codes in this table that is selected via a specified attribute value.
SPACE EFFICIENT	01h	Duplicated bytes belong to a particular object only when specific application client actions necessitate it (e.g., a copy-on-write mechanism in which duplicated bytes become associated with a particular object only when changes in their contents necessitate it).
PRE-ALLOCATED COPY ON WRITE	41h	Similar to SPACE EFFICIENT except that bytes are reserved for physical copies of all duplicated bytes (e.g., the reserved data space attribute in the User Object Information attributes page (see 7.1.3.11) is set to ensure that space is available for all duplicated bytes when application client actions necessitate copying them).
BYTE BY BYTE COPY	81h	A copy shall be made of every duplicated byte, and each object shall have its own, unique copy of the duplicated bytes.
FASTER COPY PERFORMANCE	FDh	A vendor specific object duplication mechanism whose characteristics are similar to those of the SPACE EFFICIENT object duplication method.
HIGHER DATA DUPLICATION	FEh	A vendor specific object duplication mechanism whose characteristics are similar to those of the BYTE BY BYTE COPY object duplication method.
DO NOT CARE	FFh	The device server may use any duplication method or combination of methods.
^a These codes are used in fields, attribute numbers, and attribute values. All codes not listed in this table are reserved.		

Except for the BYTE BY BYTE COPY object duplication method, the device server may use the object duplication method specified by the application client as a recommendation while still employing aspects of any supported object duplication method to achieve optimum processing times and/or space utilization.

Some object duplication methods (e.g., the SPACE EFFICIENT) cause single instances of recorded data to be shared among multiple objects. This type of shared data shall not be removed from stable storage until there are no objects that reference it.

Support for the various object duplication methods is indicated by attributes in the Root Information attributes page (see 7.1.3.8).

4.13.4 Object duplication state management

4.13.4.1 Overview

Duplicating user object data, collection data, partition data, and attributes may take a significant interval of time. The following mechanisms are provided so that application clients may specify how changes in source objects are to be handled during a duplication operation:

- a) Time of duplication (see 4.13.4.2); and
- b) Source object freeze (see 4.13.4.3).

These mechanisms are complementary. The use of one tends to eliminate the need to use any of the others.

4.13.4.2 Time of duplication source object management

Time of duplication codes (see table 44) allow the application client to specify what time in the life of a source object is to be used for purposes of duplicating that object. If multiple objects are duplicated by a single command, the time of duplication requirements apply separately to each duplicated object.

Table 44 — Time of duplication source object management

Name	Code ^a	Description
DEFAULT	0h	Used to specify one of the other codes in this table that is selected via a specified attribute value.
BEGINNING	1h	The duplicated object shall have the contents of the source object at the time the duplication was begun.
DO NOT CARE	8h	The duplicated object may have any contents of the source object, including contents that were not in effect at either the beginning or the end of the duplication.
END	Fh	The duplicated object shall have the contents of the source object at the time the duplication was completed.
^a These codes are used in field, attribute numbers, and attribute values. All codes not listed in this table are reserved.		

Support for the various time of duplication methods is indicated by attributes in the Root Information attributes page (see 7.1.3.8).

4.13.4.3 Source object freeze duplication management

A way to ensure the state of a source object during duplication is to set the object's object accessibility attribute (e.g., the object accessibility attribute in the User Object Information attributes page (see 7.1.3.11)) to disable write access (i.e., to 0000 00001h). This is described as freezing the source object.

Support for source object freeze duplication management is indicated by the support for duplicated object freezing attribute in the Root Information attributes page (see 7.1.3.8).

4.13.5 Object duplication space accounting

Some object duplication methods (e.g., the SPACE EFFICIENT object duplication method described in table 43 (see 4.13.3)) result in a situation where the used capacity (e.g., the used capacity attribute in a Partition Information attributes page (see 7.1.3.9)) of the original object plus the used capacity of the duplicate object almost

equals the used capacity of the original object alone. However, such situations also may evolve in ways where the used capacity increases for reasons that are not obvious consequences of the commands being processed. Such increases in the used capacity attribute value shall not result in CHECK CONDITION status being returned for a command that has already begun processing, but they result in quota errors being generated (see 4.10.2) for future commands.

To assist application clients in managing capacity usage and the quotas (see 4.10) on capacity usage, the potential used capacity increment attribute in the Partition Information attributes page (see 7.1.3.9) indicates the maximum number of bytes by which the used capacity attribute in the same Partition Information attributes page might increase due to ongoing command processing.

4.14 Data persistence model

The OSD data persistence model contains a two level memory hierarchy:

- a) Volatile cache – storage is:
 - A) Lost after a power on or reset event (see SAM-4); and
 - B) May be lost after an I_T nexus loss or logical unit reset event (see SAM-4);and
- b) Stable storage – storage that survives all the events that may result in the loss of data in the volatile cache.

Individual OBSD (see 3.1.27) implementations may use whatever technologies they choose to implement stable storage (e.g., an OBSD may implement stable storage as a combination of non-volatile random access memory and disk devices).

Implementation of a volatile cache is optional. Support for volatile cache, including support for the FUA bit and the DPO bit, may be indicated by setting the V_SUP bit to one in the Extended INQUIRY Data VPD page (see SPC-3).

The device server may transfer data from the volatile cache to stable storage after status has been returned for the command that placed the data in the volatile cache. Errors that occur during such data transfer operations shall be reported as deferred errors (see SPC-3).

The following bits (see 5.2.3) provide per-command controls over the use of stable storage and volatile cache:

- a) The FUA (Force Unit Access) bit controls whether or not the results of a command shall be written to stable storage before status is returned to the application client; and
- b) The DPO (Disable Page Out) bit recommends against the use of the volatile cache.

4.15 Data-In and Data-Out Buffer model

4.15.1 Bidirectional data transfers

All commands defined by this standard use both the Data-In Buffer and Data-Out Buffer.

4.15.2 OSD meta data

A single command may include the following types of data:

- a) Traditional parameter data, that is placed in a CDB continuation segment;
- b) Traditional command data;

- c) OSD object meta data; and/or
- d) Integrity check values computed over all the other types of data.

The presence of generalized object meta data differentiates communications in the OSD model from those used by traditional block structured devices (i.e., SBC devices).

NOTE 4 This standard provides for several segments in the Data-in and Data-out Buffers because the output meta data is typically too large to fit in the CDB, the input meta data is too large to fit in the single status byte returned by SCSI devices, and the ALLDATA security method (see 4.12.4.5) provides for the computation of integrity check values for all data bytes exchanged between the application client and device server.

OSD meta data and integrity check values share the Data-In Buffer and Data-Out Buffer with the traditional command or parameter data as shown in table 45.

Table 45 — OSD Data-In Buffer and Data-Out Buffer model

Bit Byte	7	6	5	4	3	2	1	0
0	Command data and/or parameter data segment, if any							
i-1								
i	Unused bytes, if any							
k-1								
k	Meta data segments, if any							
p-1								
p	Unused bytes, if any							
m-1								
m	Integrity check value segment, if any							
n								

The Data-In Buffer format is described in 4.15.3. The Data-Out Buffer format is described in 4.15.4.

Offset values (see 4.15.5) for the various segments are provided in CDB fields. The segments of the Data-In Buffer and Data-Out Buffer should not overlap. If they do, the results are unpredictable.

4.15.3 OSD Data-In Buffer format

The Data-In Buffer has the format shown in table 46.

Table 46 — OSD Data-In Buffer format

Bit Byte	7	6	5	4	3	2	1	0
0	Command data or parameter data segment, if any							
i-1								
i	Unused bytes, if any							
k-1								
k	Retrieved attributes segment, if any							
p-1								
p	Unused bytes, if any							
m-1								
m	Data-In Buffer integrity check value segment, if any (see 4.12.4.5)							
n								

The command data or parameter data segment contains data transferred from an object to the application client (e.g., data read by a READ command (see 6.27)) or data returned to the application client by the device server in response to a request made by the command (e.g., the matches list parameter data returned by a QUERY command (see 6.26)).

The retrieved attributes segment contains attribute values retrieved based on requests specified by the CDB (see 5.2.6).

The Data-In Buffer integrity check value segment contains security parameters related to the ALLDATA security method (see 4.12.4.5).

The CDB offset fields that assist in locating the Data-In Buffer segments are shown in table 47.

Table 47 — Summary of OSD Data-In Buffer offsets

CDB Data-In Buffer offset field	Reference	Buffer segment
none		Command data or parameter data
RETRIEVED ATTRIBUTES OFFSET	5.2.6	Retrieved attributes data
DATA-IN INTEGRITY CHECK VALUE OFFSET	5.2.11	Data-In Buffer integrity check value

If the device server sends data to the unused Data-In Buffer bytes in the initiator device, then the device server shall send bytes containing zero.

4.15.4 OSD Data-Out Buffer format

The Data-Out Buffer has the format shown in table 48.

Table 48 — OSD Data-Out Buffer format

Bit Byte	7	6	5	4	3	2	1	0
0	CDB continuation segment (see 5.3), if any							
h-1								
h	Command data segment, if any							
i-1								
i	Unused bytes, if any							
k-1								
k	Set attributes segment, if any							
x-1								
x	Unused bytes, if any							
y-1								
y	Get attributes segment, if any							
z-1								
z	Unused bytes, if any							
m-1								
m	Data-Out Buffer integrity check value segment, if any (see 4.12.4.5)							
n								

The CDB continuation segment contains fields that elaborate on the command to be processed (see 5.3).

The command data segment contains data to be transferred to an object from the application client (e.g., data to be written by a WRITE command (see 6.40)).

The set attributes segment contains attribute values to be set based on requests specified by the CDB (see 5.2.6).

The get attributes segment contains a list of attribute values to be retrieved based on requests specified by the CDB (see 5.2.6).

The Data-Out Buffer integrity check value segment contains security parameters related to the ALLDATA security method (see 4.12.4.5).

The CDB offset fields that assist in locating the Data-Out Buffer segments are shown in table 49.

Table 49 — Summary of OSD Data-Out Buffer offsets

CDB Data-Out Buffer offset field	Reference	Buffer segment
none		CDB continuation
CDB CONTINUATION LENGTH ^a	5.2.5	Command data
SET ATTRIBUTES LIST OFFSET	5.2.6	Set attributes
SET ATTRIBUTES OFFSET	5.2.6	Set attributes
GET ATTRIBUTES LIST OFFSET	5.2.6	Get attributes
DATA-OUT INTEGRITY CHECK VALUE OFFSET	5.2.11	Data-Out Buffer integrity check value
^a This is a count of the bytes in the preceding segment, not an offset.		

The device server shall ignore the contents of unused bytes in the Data-Out Buffer.

4.15.5 Data-In and Data-Out buffer offsets

Offset fields (see table 50) in the CDB (e.g., the RETRIEVED ATTRIBUTES OFFSET field described in 4.15.3 and the SET ATTRIBUTES LIST OFFSET field described in 4.15.4) specify the starting byte of segments in the Data-In Buffer or Data-Out Buffer other than the command data or parameter data segment.

Table 50 — CDB Data-In Buffer and Data-Out Buffer offset field format

Bit Byte	7	6	5	4	3	2	1	0
0	EXPONENT				(MSB)			
1								
2								
3	MANTISSA							
	(LSB)							

The EXPONENT field specifies the signed integer (see 3.6) power of two to be used in computing the byte offset. The power of two shall be the value in the EXPONENT field plus eight. If the OFFSET field does not contain FFFF FFFFh and the EXPONENT field contains -6, -7, or -8, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The MANTISSA field specifies the value to be multiplied by two raised to the power specified by the EXPONENT field.

The byte offset represented by a field having the format described in this subclause shall be:

$$\text{byte offset} = \text{mantissa} \times (2^{(\text{exponent}+8)})$$

An offset field containing zero specifies a byte offset value of zero.

If the offset field for a Data-In Buffer or Data-Out Buffer segment that is not being used is not set to FFFF FFFFh, the results are unpredictable.

4.16 Error reporting

4.16.1 Introduction

OSD logical units shall use descriptor format sense data (see SPC-3) to report all errors.

All sense data returned by OSD device servers shall include the OSD error identification sense data descriptor (see 4.16.2.1) to identify the OSD object in which the reported error was detected.

If it is possible to identify a specific byte or range of bytes within a user object as being associated with an error, the information sense data descriptor (see SPC-3) shall be included in the sense data with the INFORMATION field set to the byte within the user object associated with the error or the first byte in the range of bytes within the user object associated with the error.

If a READ command (see 6.27) attempts to read bytes both before and beyond a user object's logical length, the command-specific information sense data descriptor (see SPC-3) shall be included in the sense data with the COMMAND-SPECIFIC INFORMATION field set to the number of bytes transferred before the user object's logical length was reached.

If the CMDRSP security method or the ALLDATA security method (see 4.12.4) is used to process the command, the sense data shall include the OSD response integrity check value sense data descriptor (see 4.16.2.2). If the status is not CHECK CONDITION and no sense data is transferred, the response integrity check value is returned in the response integrity check value attribute in the Current Command attributes page (see 7.1.3.31).

The OSD CDB is very large. To reduce uncertainty in determining errors in CDB field settings or in parameter data, any sense data having the sense key set to ILLEGAL REQUEST should include the sense key specific sense data descriptor (see SPC-3) with the field pointer sense key specific data.

Errors other than those defined in this standard may be reported as needed. The sense data shall include the appropriate sense key and additional sense code (see SPC-3) to identify the condition.

Errors may occur after the command has completed. For such errors, SPC-3 defines a deferred error reporting mechanism.

4.16.2 OSD-specific sense data descriptors

4.16.2.1 OSD error identification sense data descriptor

The OSD object identification sense data descriptor (see table 51) provides information that identifies the OSD object associated with the error reported in the sense data (see OSD).

Table 51 — OSD object identification sense data descriptor format

Bit Byte	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE (06h)							
1	ADDITIONAL LENGTH (1Eh)							
2	Reserved							
7								
8	NOT INITIATED COMMAND FUNCTIONS							
11								
12	COMPLETED COMMAND FUNCTIONS							
15								
16	(MSB)	PARTITION_ID						(LSB)
23								
24	(MSB)	OBJECT_ID						(LSB)
31								

The NOT INITIATED COMMAND FUNCTIONS field contains the command functions bits (see table 52) that indicate (see table 54) which command functions had not been initiated at the time the error reported in the sense data was detected.

The COMPLETED COMMAND FUNCTIONS field contains the command functions bits (see table 52) that indicate (see table 54) which command functions had been completed at the time the error reported in the sense data was detected.

The PARTITION_ID field contains the Partition_ID (see 4.6.4) of the partition that is associated with the error being reported.

The OBJECT_ID field contains the Collection_Object_ID (see 4.6.6) or User_Object_ID (see 4.6.5) of the object that is associated with the error being reported.

The command functions bits (see table 52) are contained in the NOT INITIATED COMMAND FUNCTIONS field and COMPLETED COMMAND FUNCTIONS field (see table 51).

Table 52 — Command functions bits

Bit Byte	7	6	5	4	3	2	1	0
0	VALIDATION	Reserved	CMD_CAP_V	COMMAND	Reserved	Reserved	Reserved	Reserved
1	Reserved	Reserved	Reserved	IMP_ST_ATT	Reserved	Reserved	Reserved	Reserved
2	Reserved	Reserved	SA_CAP_V	SET_ATT	Reserved	Reserved	Reserved	Reserved
3	Reserved	Reserved	GA_CAP_V	GET_ATT	Reserved	Reserved	Reserved	Reserved

The command functions bits and the command functions that they indicate are listed in table 53.

Table 53 — Command functions indicated by the command functions bits

Command functions bit	Command function indicated
VALIDATION	Validation of the command, including security parameters
CMD_CAP_V	Capability verification for those command functions not related to attributes (e.g., writing data to a user object)
COMMAND	Processing of those command functions not related to attributes
IMP_ST_ATT	Processing of any set attributes command functions resulting from the processing of the command (e.g., changes due to a WRITE command)
SA_CAP_V	Capability verification for all set attributes command functions specified in the CDB
SET_ATT	Processing of any set attributes command functions specified in the CDB
GA_CAP_V	Capability verification for all get attributes command functions specified in the CDB
GET_ATT	Processing of any get attributes command functions specified in the CDB

The interpretation of the combinations of the command functions bits in the NOT INITIATED COMMAND FUNCTIONS field and COMPLETED COMMAND FUNCTIONS field is shown in table 54.

Table 54 — Command functions bits combinations

NOT INITIATED COMMAND FUNCTIONS bit	COMPLETED COMMAND FUNCTIONS bit	Status of the indicated command function at the time the error reported by the sense data was detected
1	0	Processing was requested, but was not initiated and not completed
0	0	Processing was not requested, or processing was in progress
0	1	Processing was requested and completed
1	1	Reserved

4.16.2.2 OSD response integrity check value sense data descriptor

The OSD response integrity check value sense data descriptor (see table 55) contains the response integrity check value used when the OSD security method is CMDRSP or ALLDATA (see 4.12.4).

Table 55 — OSD response integrity check value sense data descriptor format

Bit Byte	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE (07h)							
1	ADDITIONAL LENGTH (20h)							
2	(MSB)							
33	RESPONSE INTEGRITY CHECK VALUE (LSB)							

The RESPONSE INTEGRITY CHECK VALUE field contains the response integrity check value (see 4.12.8) that is computed as described in 4.12.4.4 for the command for which the error being reported.

4.16.2.3 OSD attribute identification sense data descriptor

The OSD attribute identification sense data descriptor (see table 56) identifies one or more attributes (see 7.1) associated with the error reported in the sense data.

Table 56 — OSD attribute identification sense data descriptor format

Bit Byte	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE (08h)							
1	ADDITIONAL LENGTH (n-2)							
2	Reserved							
3	Reserved							
	Attribute descriptors							
4	Attribute descriptor 0 (see table 57)							
	⋮							
n	Attribute descriptor x (see table 57)							

Each attribute descriptor (see table 57) identifies one attribute associated with the error reported in the sense data.

Table 57 — Sense data attribute descriptor format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							(LSB)
3	ATTRIBUTE PAGE							(LSB)
4	(MSB) _____							(LSB)
7	ATTRIBUTE NUMBER							(LSB)

The ATTRIBUTE PAGE field contains the attribute page number (see 4.8.5) for the attributes page containing the attribute associated with the error reported in the sense data.

The ATTRIBUTE NUMBER field contains the attribute number (see 4.8.6) of the attribute associated with the error reported in the sense data.

4.16.3 Auto contingent allegiance

OSD logical units that are not capable of accepting a command with the ACA task attribute (see SAM-3) at any time and performing all data transfer operations that the command requests shall set the NORMACA bit to zero in the Standard INQUIRY data (see SPC-3).

4.17 Reservations

The access enabled or access disabled condition determines when an application client may store or retrieve user data on all or part of the medium. Access may be restricted for read command functions, write command functions, or both. This attribute may be controlled by an external mechanism or by the PERSISTENT RESERVE IN command and PERSISTENT RESERVE OUT command (see SPC-3). The OSD logical unit shall not support the RESERVE command or the RELEASE command.

The credential-based system defined by the OSD security model (see 4.12) provides access controls that are more appropriate to an OBSD (see 3.1.27) than persistent reservations. Use of persistent reservations is permitted only if use of the OSD security model is not activated. If the security method in effect for the root or any partition in the OSD logical unit is not NOSEC (see 4.12.4), the PERSISTENT RESERVE IN command and PERSISTENT RESERVE OUT command shall be treated as invalid commands (see SPC-3).

If a persistent reservation is in effect or any registrations are established when the security method in effect for the root or any partition changes from the NOSEC security method to any other security method, the persistent reservation, if any, shall be released and all registrations shall be unregistered. A unit attention condition (see SAM-3) shall be established for the initiator port associated with every registered I_T nexus. The sense key shall be set to UNIT ATTENTION and the additional sense code shall be set to RESERVATIONS RELEASED.

The PERSISTENT RESERVE IN command and the PERSISTENT RESERVE OUT command define how different types of restricted access may be achieved, and to whom the access is restricted. This subclause describes the interaction of the application client that requested the reservation, and the other application clients.

An application client uses reservations to gain a level of exclusivity in access to all or part of the medium for itself or another application client. It is expected that the reservation is retained until released. The device server ensures

that the application client with the reservation is able to access the reserved media within the operating parameters established by that application client.

Reservation restrictions are placed on commands as a result of access qualifiers associated with the type of reservation. The details of commands that are allowed under what types of reservations are described in table 58.

Commands from initiator ports holding a reservation should complete normally. The behavior of commands from registered initiator ports when a registrants only or all registrants persistent reservation is present is specified in table 58.

A command shall be checked for reservation conflicts before the task containing that command enters the enabled task state.

For each command, this standard or SPC-3 defines the conditions that result in RESERVATION CONFLICT.

Table 58 — OSD commands that are allowed in the presence of various reservations (part 1 of 2)

OSD Command	Addressed logical unit has this type of persistent reservation held by another I_T nexus				
	From any I_T nexus		From registered I_T nexus (RR all types)	From not registered I_T nexus	
	Write Excl	Excl Access		Write Excl RR	Excl Access – RR
APPEND	Conflict	Conflict	Allowed	Conflict	Conflict
CLEAR	Conflict	Conflict	Allowed	Conflict	Conflict
COPY USER OBJECTS	Conflict	Conflict	Allowed	Conflict	Conflict
CREATE	Conflict	Conflict	Allowed	Conflict	Conflict
CREATE AND WRITE	Conflict	Conflict	Allowed	Conflict	Conflict
CREATE COLLECTION	Conflict	Conflict	Allowed	Conflict	Conflict
CREATE CLONE	Conflict	Conflict	Allowed	Conflict	Conflict
CREATE PARTITION	Conflict	Conflict	Allowed	Conflict	Conflict
CREATE SNAPSHOT	Conflict	Conflict	Allowed	Conflict	Conflict
CREATE USER TRACKING COLLECTION	Conflict	Conflict	Allowed	Conflict	Conflict
DETACH CLONE	Conflict	Conflict	Allowed	Conflict	Conflict
FLUSH	Conflict	Conflict	Allowed	Conflict	Conflict
FLUSH COLLECTION	Conflict	Conflict	Allowed	Conflict	Conflict
FLUSH OSD	Conflict	Conflict	Allowed	Conflict	Conflict
FLUSH PARTITION	Conflict	Conflict	Allowed	Conflict	Conflict
FORMAT OSD	Conflict	Conflict	Allowed	Conflict	Conflict
GET ATTRIBUTES	Allowed	Conflict	Allowed	Allowed	Conflict
GET MEMBER ATTRIBUTES	Allowed	Conflict	Allowed	Allowed	Conflict
LIST	Allowed	Conflict	Allowed	Allowed	Conflict
Key: Excl =Exclusive, RR =Registrants Only or All Registrants					

Table 58 — OSD commands that are allowed in the presence of various reservations (part 2 of 2)

OSD Command	Addressed logical unit has this type of persistent reservation held by another I_T nexus				
	From any I_T nexus		From registered I_T nexus (RR all types)	From not registered I_T nexus	
	Write Excl	Excl Access		Write Excl RR	Excl Access – RR
LIST COLLECTION	Allowed	Conflict	Allowed	Allowed	Conflict
OBJECT STRUCTURE CHECK	Conflict	Conflict	Allowed	Conflict	Conflict
PERFORM SCSI COMMAND	Conflict	Conflict	Allowed	Conflict	Conflict
PERFORM TASK MANAGEMENT FUNCTION	Conflict	Conflict	Allowed	Conflict	Conflict
PUNCH	Conflict	Conflict	Allowed	Conflict	Conflict
QUERY	Allowed	Conflict	Allowed	Allowed	Conflict
READ	Allowed	Conflict	Allowed	Allowed	Conflict
READ MAP	Conflict	Conflict	Allowed	Conflict	Conflict
READ MAPS AND COMPARE	Conflict	Conflict	Allowed	Conflict	Conflict
REFRESH SNAPSHOT OR CLONE	Conflict	Conflict	Allowed	Conflict	Conflict
REMOVE	Conflict	Conflict	Allowed	Conflict	Conflict
REMOVE COLLECTION	Conflict	Conflict	Allowed	Conflict	Conflict
REMOVE MEMBER OBJECTS	Conflict	Conflict	Allowed	Conflict	Conflict
REMOVE PARTITION	Conflict	Conflict	Allowed	Conflict	Conflict
RESTORE PARTITION FROM SNAPSHOT	Conflict	Conflict	Allowed	Conflict	Conflict
SET ATTRIBUTES	Conflict	Conflict	Allowed	Conflict	Conflict
SET KEY	Conflict	Conflict	Allowed	Conflict	Conflict
SET MASTER KEY	Conflict	Conflict	Allowed	Conflict	Conflict
SET MEMBER ATTRIBUTES	Conflict	Conflict	Allowed	Conflict	Conflict
WRITE	Conflict	Conflict	Allowed	Conflict	Conflict
Any command that retrieves attributes	see the command entry in this table				
Any command that sets attributes	Conflict	Conflict	Allowed	Conflict	Conflict
Key: Excl =Exclusive, RR =Registrants Only or All Registrants					

5 Common Formats

5.1 OSD CDB format

The OSD CDB consists of a 10-byte header followed by service action specific fields.

An application client sends a CDB to the device server. If a device server receives a CDB containing an operation code that is invalid or not supported, it shall return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and an additional sense code set to INVALID COMMAND OPERATION CODE. If a device server receives a CDB containing a service action that is invalid or not supported, it shall return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and an additional sense code set to INVALID FIELD IN CDB.

The OSD commands defined in this standard use the variable length CDB format (see SPC-3). In the variable length CDB (see table 59), an OPERATION CODE field containing 7Fh is the first byte and a CONTROL byte is the second byte. The general structure of the OPERATION CODE field and CONTROL byte are defined in SAM-3.

Table 59 — Basic OSD CDB

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (7Fh)							
1	CONTROL							
2	Reserved							
3	Reserved							
4	Reserved							
5	Reserved							
6	Reserved							
7	ADDITIONAL CDB LENGTH (228)							
8	(MSB)							
9	SERVICE ACTION							
10	(LSB)							
235	Service action specific fields (see 5.2.1)							

The ADDITIONAL CDB LENGTH field specifies the number of bytes following it in the variable length CDB. If the value in the ADDITIONAL CDB LENGTH field is not 216, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The SERVICE ACTION field indicates the action being requested by the application client. Each service action code description defines the service action specific fields that are needed for that service action.

5.2 Fields commonly used in OSD commands

5.2.1 Overview

OSD commands employ the basic CDB structure shown in 5.1. Within the basic CDB structure, the OSD service action specific fields are organized so that the same field is in the same location in all OSD CDBs (see table 60). OSD service action specific fields that are unique to a small number of CDBs are not shown in this subclause.

Table 60 — OSD service action specific fields

Bit Byte	7	6	5	4	3	2	1	0
10	Reserved			DPO ^a	FUA ^a	ISOLATION (see 5.2.8)		
11	IMMED_TR ^b	Reserved	GET/SET CDBFMT ^c		Command specific options			
12	TIMESTAMPS CONTROL (see 5.2.13)							
13	Reserved							
15								
16	(MSB)	PARTITION_ID (see 5.2.10)						(LSB)
23								
24	(MSB)	USER_OBJECT_ID (see 5.2.14)						(LSB)
31								
32	(MSB)	LENGTH (see 5.2.9) or ALLOCATION LENGTH (see 5.2.2)						(LSB)
39								
40	(MSB)	STARTING BYTE ADDRESS (see 5.2.12)						(LSB)
47								
48	(MSB)	CDB CONTINUATION LENGTH (see 5.2.5)						(LSB)
51								
52	Get and set attributes parameters ^c							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								
^a See 5.2.3. ^b See 5.2.5. ^c See 5.2.6.								

5.2.2 Allocation length

The ALLOCATION LENGTH field specifies the maximum number of bytes that an application client has allocated for the return of parameter data. An allocation length of zero indicates that no data shall be transferred. This condition shall not be considered as an error.

The allocation length is used to limit the maximum amount of the parameter data returned to an application client. The device server shall terminate transfers to the Data-In Buffer if the number of bytes specified by the ALLOCATION LENGTH field have been transferred or if all available data have been transferred, whichever is less. If the information being transferred is truncated, the contents of the parameter data ADDITIONAL LENGTH field shall not be altered to reflect the truncation.

5.2.3 Caching control bits

The DPO (disable page out) bit allows the application client to influence the use of volatile cache (see 4.14). If the DPO bit is set to zero, the use of volatile cache should proceed without influence caused by the DPO bit value. If the DPO bit is set to one, the device server should not place data transferred as a result of this command in the volatile cache.

The FUA (force unit access) bit controls whether or not the results of a command shall be written to stable storage (see 4.14) before status is returned to the application client. If the FUA bit is set to zero, the device server may return status as soon as the data transferred by this command is in the volatile cache. If the FUA bit is set to one, the device server shall not return status until the data transferred by this command (i.e., either read data or write data) has been written to stable storage.

The direction of data transfer has no effect on the meaning of the DPO and FUA bits. The DPO and FUA bits affect the processing of both OSD object data and attributes.

5.2.4 Capability

The capability is described in 4.11.2.2. Any security method other than NOSEC (see 4.12.4) is design to return an error unless the capability in the CDB is a solo capability (see 4.12.5).

5.2.5 CDB continuation length

The CDB CONTINUATION LENGTH field specifies the number of bytes in the CDB continuation segment of the Data-Out Buffer (see 4.15.4) using the format described in 5.3.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB, if the CDB CONTINUATION LENGTH field contains a non-zero value that is:

- a) Not a multiple of eight;
- b) Less than 48 (i.e., 30h); or
- c) Greater than the value in the maximum CDB continuation length attribute in the Root Information attributes page (see 7.1.3.8).

5.2.6 Get and set attributes parameters

5.2.6.1 Get and set attributes CDB format selection

The GET/SET CDBFMT (get and set attributes CDB format) field (see table 61) specifies the format of the get and set attributes parameters in the CDB.

Table 61 — Get and set attributes CDB format code values

Value	Description	Reference
00b	Reserved	
01b	Set one attribute using CDB fields	5.2.6.2
10b	Get an attributes page and set an attribute value	5.2.6.3
11b	Get and set attributes using lists	5.2.6.4

5.2.6.2 Set one attribute value using CDB fields

The set one attribute using CDB parameters format (see table 62) allows the setting of a single attribute using CDB fields and does not allow the retrieval of any attributes.

Table 62 — Set one attribute using CDB parameters format

Bit Byte	7	6	5	4	3	2	1	0
	⋮ Other CDB fields							
47								
48	(MSB)	ATTRIBUTES PAGE						(LSB)
51								
52	(MSB)	ATTRIBUTE NUMBER						(LSB)
55								
56	(MSB)	ATTRIBUTE LENGTH						(LSB)
57								
58	(MSB)	ATTRIBUTE VALUE						(LSB)
75								
76	⋮ Other CDB fields							

The ATTRIBUTES PAGE field specifies the page number of the attribute value to be set.

The ATTRIBUTE NUMBER field specifies the attribute number within the attributes page specified by the ATTRIBUTES PAGE field of the attribute value to be set.

If the ATTRIBUTES PAGE field or ATTRIBUTE NUMBER field contains FFFF FFFFh, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The ATTRIBUTE LENGTH field specifies the length of the attribute value in bytes. If the attribute length is greater than 18, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the ATTRIBUTE LENGTH field does not contain zero, the ATTRIBUTE VALUE field specifies the attribute value. If the ATTRIBUTE LENGTH field contains zero, the contents of the ATTRIBUTE VALUE field are ignored. The attribute length field shall affect the setting of the attribute as described in 4.8.3.

If setting an attribute value causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the capacity quota.

5.2.6.3 Get an attributes page and set an attribute value

The page oriented get and set attributes parameters CDB format (see table 63) allows the retrieval of one attributes page and the setting of one attribute value.

Table 63 — Page oriented get and set attributes CDB parameters format

Bit Byte	7	6	5	4	3	2	1	0
	⋮ Other CDB fields							
47								
48	(MSB)	GET ATTRIBUTES PAGE						(LSB)
51								
52	(MSB)	GET ATTRIBUTES ALLOCATION LENGTH						(LSB)
55								
56								
59	RETRIEVED ATTRIBUTES OFFSET							
60	(MSB)	SET ATTRIBUTES PAGE						(LSB)
63								
64	(MSB)	SET ATTRIBUTE NUMBER						(LSB)
67								
68	(MSB)	SET ATTRIBUTE LENGTH						(LSB)
71								
72								
75	SET ATTRIBUTES OFFSET							
76	⋮ Other CDB fields							

The GET ATTRIBUTES PAGE field specifies the attributes page number (see 7.1.3) to be retrieved. Zero specifies that no attributes page is to be retrieved.

The GET ATTRIBUTES ALLOCATION LENGTH field specifies the number of bytes allocated to receive the retrieved attributes page.

If the get attributes allocation length is not sufficient to accommodate all bytes in the specified attributes page, the transfer of attributes data shall be truncated at the specified get attributes allocation length and this shall not be considered to be an error. If get attributes data is truncated, all the get attributes data that is transferred shall be transferred as if no error occurred (i.e., length fields in the transferred get attributes data shall not be modified to reflect the truncation).

The RETRIEVED ATTRIBUTES OFFSET field specifies the byte offset of the first Data-In Buffer byte to contain the retrieved attributes page. The format of the RETRIEVED ATTRIBUTES OFFSET field is described in 4.15.5. The format of the Data-In Buffer when attributes are being retrieved is described in 4.15.3.

The SET ATTRIBUTES PAGE field and SET ATTRIBUTE NUMBER field specify one attribute value to be set. A zero in the SET ATTRIBUTES PAGE field specifies that no attribute value is to be set.

The SET ATTRIBUTE LENGTH field specifies the number of bytes in the attribute being set.

The SET ATTRIBUTES OFFSET field specifies the byte offset of the first Data-Out Buffer byte containing the value of the attribute to be set. The format of the SET ATTRIBUTES OFFSET field is described in 4.15.5. The format of the Data-Out Buffer when attributes are being set is described in 4.15.4.

If the set attributes page is non-zero and the attribute specified by the SET ATTRIBUTES PAGE field and SET ATTRIBUTE NUMBER field is application client settable, the attribute length shall be set to the value in the SET ATTRIBUTE LENGTH field and the value of the attribute shall be set to the contents of the Data-Out Buffer starting at the byte offset specified by the SET ATTRIBUTES OFFSET field and continuing for the number of bytes specified by the SET ATTRIBUTE LENGTH field. If the attribute specified by the SET ATTRIBUTES PAGE field and SET ATTRIBUTE NUMBER field has not been defined previously setting it shall not be considered an error.

If setting an attribute value causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the capacity quota.

5.2.6.4 Get and set attributes lists

The list oriented get and set attributes parameters CDB format (see table 64) allows the retrieval and setting of attributes using lists (see 7.1.4).

Table 64 — List oriented get and set attributes CDB parameters format

Bit Byte	7	6	5	4	3	2	1	0
	⋮ Other CDB fields							
47								
48	(MSB)	GET ATTRIBUTES LIST LENGTH						(LSB)
51								
52	GET ATTRIBUTES LIST OFFSET							
55								
56	(MSB)	GET ATTRIBUTES ALLOCATION LENGTH						(LSB)
59								
60	RETRIEVED ATTRIBUTES OFFSET							
63								
64	(MSB)	SET ATTRIBUTES LIST LENGTH						(LSB)
67								
68	SET ATTRIBUTES LIST OFFSET							
71								
72	Reserved							
75								
76	⋮ Other CDB fields							

The GET ATTRIBUTES LIST LENGTH field specifies the length of a get attributes list (see 7.1.4) that specifies one or more attribute values to be retrieved. A get attributes list length of zero specifies that no get attributes list is included with the command.

The GET ATTRIBUTES LIST OFFSET field specifies the byte offset of the first Data-Out Buffer byte containing the get attributes list. The format of the GET ATTRIBUTES LIST OFFSET field is described in 4.15.5. The format of the Data-Out Buffer when a list is being use to retrieve attributes is described in 4.15.4.

The GET ATTRIBUTES ALLOCATION LENGTH field specifies the number of bytes allocated to receive retrieved attributes page.

If the get attributes allocation length is not sufficient to accommodate all bytes in the attributes specified by the get attributes list, the transfer of attributes data shall be truncated at the specified get attributes allocation length, this shall not be considered to be an error. If get attributes data is truncated, all the get attributes data that is transferred shall be transferred as if no error occurred (i.e., length fields in the transferred get attributes data shall not be modified to reflect the truncation).

The RETRIEVED ATTRIBUTES OFFSET field specifies the byte offset of the first Data-In Buffer byte to contain the retrieved attributes list. The format of the RETRIEVED ATTRIBUTES OFFSET field is described in 4.15.5. The format of the Data-In Buffer when attributes are being retrieved is described in 4.15.3.

The SET ATTRIBUTES LIST LENGTH field specifies the length of a set attributes list (see 7.1.4) that specifies one or more attribute values to be set. A set attributes list length of zero specifies that there is no set attributes list included with the command.

If the set attributes parameter list length causes the truncation of the attribute value or entry in the set attributes list, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The SET ATTRIBUTES LIST OFFSET field specifies the byte offset of the first Data-Out Buffer byte containing the first byte of the set attributes list. The format of the SET ATTRIBUTES OFFSET field is described in 4.15.5. The format of the Data-Out Buffer when attributes are being set is described in 4.15.4.

If setting an attribute value causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the capacity quota.

5.2.7 Immediate bit for TRACKING collections

The immediate bit for TRACKING collections (IMMED_TR) allows an application client to specify that the command be completed with GOOD status after the TRACKING collection (see 4.6.6.3) has been set up but before all objects in the TRACKING collection have been processed. If the IMMED_TR bit is set to zero, the device server shall process all command functions (e.g., all objects in the TRACKING collection evaluated) before completing the command. If the IMMED_TR bit is set to one, the device server shall:

- 1) Verify the correctness of all CDB and CDB continuation fields;
- 2) Perform any security checks required to validate the command (see 4.12);
- 3) Initialize the Command Tracking attributes page (see 7.1.3.20) in the TRACKING collection with all information necessary to process, track, and restart the command;
- 4) Process all command functions related to attributes as described in 4.8.4, except those command functions that are to be performed individually on objects in the TRACKING collection; and
- 5) If no errors have been detected, complete the command with GOOD status.

Commands completed with GOOD status due to the IMMED_TR bit being set shall not result in the REQUEST SENSE command reporting progress indication information (see SPC-4).

5.2.8 Isolation

The ISOLATION field (see table 65) specifies the isolation mode to be applied to this command.

Table 65 — ISOLATION field

Code	Description
0h	The isolation method specified by the default isolation method attribute in the Root Information attributes page (see 7.1.3.8) shall be applied to this command.
1h	The NONE isolation method described in table 162 (see 7.1.3.8) shall be applied to this command.
2h	The STRICT isolation method described in table 162 shall be applied to this command.
3h	Reserved
4h	The RANGE isolation method described in table 162 shall be applied to this command.
5h	The FUNCTIONAL isolation method described in table 162 shall be applied to this command.
6h	Reserved
7h	Vendor specific

If the ISOLATION field contains a value that the supported isolation methods attribute in the Root Information attributes page (see 7.1.3.8) indicates is not supported, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

5.2.9 Length

The LENGTH field specifies the number of bytes to be transferred by a read or write.

The format of the Data-In Buffer and Data-Out Buffer is described in 4.15

5.2.10 Partition_ID

The PARTITION_ID field contains the Partition_ID (see 4.6.4) that the command is to act upon. If the partition identified by the PARTITION_ID field does not exist, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

5.2.11 Security parameters

The CDB security parameters (see table 66) contain the security information needed for each command.

Table 66 — Security parameters format

Bit Byte	7	6	5	4	3	2	1	0
	⋮ Other CDB fields							
183								
184	(MSB)							
215	REQUEST INTEGRITY CHECK VALUE (LSB)							
216								
227	REQUEST NONCE							
228								
231	DATA-IN INTEGRITY CHECK VALUE OFFSET							
232								
235	DATA-OUT INTEGRITY CHECK VALUE OFFSET							

The REQUEST INTEGRITY CHECK VALUE field contains an integrity check value (see 4.12.8) for the request sent by the application client. The REQUEST INTEGRITY CHECK VALUE field is used only by the CAPKEY security method, the CMDRSP security method, and the ALLDATA security method (see 4.12.4).

The CAPKEY security method for computing the request integrity check value is described in 4.12.4.3. The CMDRSP security method and ALLDATA security method for computing the request integrity check value is described in 4.12.4.4.

The device server shall validate the request integrity check value as described in 4.12.6.1.

For the CMDRSP security method and the ALLDATA security method (see 4.12.4), the REQUEST NONCE field contains a security nonce (see 4.12.7). Otherwise, the REQUEST NONCE field should contain zero.

The device server shall validate the request nonce as described in 4.12.7.2 and 4.12.6.1.

The DATA-IN INTEGRITY CHECK VALUE OFFSET field specifies the byte offset of the first Data-In Buffer byte containing integrity check value information for the Data-In Buffer. If the command is not prepared for processing using the ALLDATA security method (see 4.12.4), the DATA-IN INTEGRITY CHECK VALUE OFFSET field contains FFFF FFFFh. Otherwise, the DATA-IN INTEGRITY CHECK VALUE OFFSET field contains an offset value (see 4.15.5) that specifies the first byte of the data-in integrity information that is prepared and validated as described in 4.12.4.5. The format of the Data-In Buffer when the data-in integrity check information is present is described in 4.15.3.

The DATA-OUT INTEGRITY CHECK VALUE OFFSET field specifies the byte offset of the first Data-Out Buffer byte containing integrity check value information for the Data-Out Buffer. If the command is not prepared for processing using the ALLDATA security method, the DATA-OUT INTEGRITY CHECK VALUE OFFSET field contains FFFF FFFFh. Otherwise, the DATA-OUT INTEGRITY CHECK VALUE OFFSET field contains an offset value (see 4.15.5) that specifies the first byte of the data-out integrity information that is prepared and validated as described in 4.12.4.5. The format of the Data-Out Buffer when the data-out integrity check information is present is described in 4.15.4.

5.2.12 Starting byte address

The STARTING BYTE ADDRESS field specifies the location where the read or write is to commence in the specified object relative to the first byte (i.e., byte zero) of the user object.

The format of the Data-In Buffer and Data-Out Buffer is described in 4.15

5.2.13 Timestamps control

The TIMESTAMPS CONTROL field specifies the timestamp update policy (see table 67) if the following conditions are met:

- a) If the bypass timestamps attribute in the Root Timestamps attributes page (see 7.1.3.15) contains FFh and the command is:
 - A) A CREATE PARTITION command;
 - B) A FLUSH OSD command;
 - C) A FORMAT OSD command;
 - D) A GET ATTRIBUTES command addressed to the root object;
 - E) A LIST command addressed to the root object;
 - F) A PERFORM SCSI COMMAND command addressed to the root object;
 - G) A PERFORM TASK MANAGEMENT FUNCTION command addressed to the root object;
 - H) A REMOVE PARTITION command;
 - I) A SET ATTRIBUTES command addressed to the root object;
 - J) A SET KEY command with the KEY TO SET field set to 01b; or
 - K) A SET MASTER KEY command;
 or
- b) If the command is not one of those listed in item a) and bypass timestamps attribute in the Partition Timestamps attributes page (see 7.1.3.16) contains FFh.

Table 67 — Timestamps control values

Value	Description
0h	Timestamps shall updated as described in the subclause that defines them
01h to 7Eh	Reserved
7Fh	Timestamps shall not be updated
80h to DFh	Reserved
E0h to FFh	Vendor specific

A timestamp attribute (see 7.1) that has never been updated shall have a length of six and a value of zero.

Bypassing a timestamp update shall not affect any previously established timestamp attribute values.

5.2.14 User_Object_ID

The USER_OBJECT_ID field contains the User_Object_ID of the user object (see 4.6.5) upon which the command is to act. If the user object identified by the USER_OBJECT_ID field does not exist, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

5.3 CDB continuation segment format

If the CDB CONTINUATION LENGTH field (see 5.2.1) is not set to zero, the first bytes in the Data-Out Buffer (see 4.15.4) have the format shown in table 68.

Table 68 — CDB continuation segment format

Bit Byte	7	6	5	4	3	2	1	0						
0	CDB CONTINUATION FORMAT													
1	Reserved													
2	(MSB)	CONTINUED SERVICE ACTION						(LSB)						
3														
4														
7														
8	(MSB)	CONTINUATION INTEGRITY CHECK VALUE						(LSB)						
39														
CDB continuation descriptors														
40	CDB continuation descriptor [first] (see 5.4.1)													
⋮														
	CDB continuation descriptor [last] (see 5.4.1)													
n	Pad bytes (for application client alignment)													

The CDB CONTINUATION FORMAT field (see table 69) specifies the format of this CDB continuation segment.

Table 69 — CDB CONTINUATION FORMAT field

Value	Description
00h	Reserved
01h	The format defined by this standard
02h to FFh	Reserved

The CONTINUED SERVICE ACTION field specifies the service action for the command to which the CDB continuation is being applied. If the contents of the CONTINUED SERVICE ACTION field do not match the contents of the SERVICE ACTION field in the CDB (see 5.1), then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The CONTINUATION INTEGRITY CHECK VALUE field contains an integrity check value (see 4.12.8) for the CDB continuation segment sent by the application client. The CONTINUATION INTEGRITY CHECK VALUE field is used for all security methods (see 4.12.4) except NOSEC.

Each CDB continuation descriptor (see 5.4) contains one set of CDB continuation information. Unless otherwise stated, the order in which the CDB continuation descriptors appear in the CDB continuation segment has no significance.

The CDB continuation segment may be padded to meet alignment requirements determined by the application client. Depending on the needed alignment, zero or more bytes containing zeros may be added at the end of the CDB continuation segment.

5.4 CDB continuation descriptors

5.4.1 Overview

Each CDB continuation descriptor (see table 70) contains one set of CDB continuation information formatted as a header with a format that is common to all CDB continuation descriptors followed by data that is specific to each CDB continuation descriptor type.

Table 70 — CDB continuation descriptor format

Bit Byte	7	6	5	4	3	2	1	0
	CDB continuation descriptor header							
0	(MSB)	CDB CONTINUATION DESCRIPTOR TYPE						(LSB)
1								
2	Reserved							
3	Reserved				PAD LENGTH (p-n)			
4	(MSB)	CDB CONTINUATION DESCRIPTOR LENGTH (n-7)						(LSB)
7								
	CDB continuation descriptor type specific data							
8	CDB continuation descriptor type specific data							
n								
	CDB continuation descriptor alignment bytes							
n+1	Pad bytes (for eight-byte alignment)							
p								

The CDB CONTINUATION DESCRIPTOR TYPE field (see table 71) specifies the format of the CDB continuation descriptor type specific data.

Table 71 — CDB CONTINUATION DESCRIPTOR TYPE field

Value	Description	Reference
0000h	No more continuation descriptors ^a	
0001h	Scatter/gather list	5.4.2
0002h	Query list	5.4.3
0100h	User object	5.4.4
0101h	Copy user object source	5.4.5
FFEEh	Extension capabilities	5.4.6
all other values	Reserved	
^a Since the CDB continuation segment pad bytes, if any, are set to zero (see 5.3), encountering a CDB continuation descriptor type of zero shall be processed in the same way as reaching the last byte of the CDB continuation segment.		

The PAD LENGTH field specifies the number of bytes containing zeros that follow the CDB continuation descriptor type specific data.

The CDB CONTINUATION DESCRIPTOR LENGTH field contains the number of bytes of CDB continuation descriptor type specific data that follow in this descriptor. If the value in the CDB CONTINUATION DESCRIPTOR LENGTH field is greater than the value of the supported CDB continuation descriptor type attribute for the CDB continuation descriptor type in the Root Information attributes page (see 7.1.3.8) for the CDB continuation descriptor type, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

If the sum of the pad length and the CDB continuation descriptor length is not a multiple of eight, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The format of the CDB continuation descriptor type specific data depends on the contents of the CDB CONTINUATION DESCRIPTOR TYPE field (see table 71).

5.4.2 Scatter/gather list

The scatter/gather list CDB continuation descriptor (see table 72) specifies the relationship between contents of the command data buffer segment (see 4.15.3 for the Data-In Buffer or 4.15.4 for the Data-Out Buffer) and the bytes in the user object.

Table 72 — Scatter/gather list CDB continuation descriptor format

Bit Byte	7	6	5	4	3	2	1	0
	CDB continuation descriptor header							
0	(MSB) _____ CDB CONTINUATION DESCRIPTOR TYPE (0001h) _____ (LSB)							
1								
2	Reserved							
3	Reserved				PAD LENGTH (000b)			
4	(MSB) _____ CDB CONTINUATION DESCRIPTOR LENGTH (n-7) _____ (LSB)							
7								
	CDB continuation descriptor type specific data							
8	Scatter/gather list entry [first] (see table 73)							
23								
	⋮							
n-15	Scatter/gather list entry [last] (see table 73)							
n								

The CDB CONTINUATION DESCRIPTOR TYPE field contains 0001h (i.e., scatter/gather list CDB continuation descriptor).

The PAD LENGTH field is set to zero to indicate that no pad bytes are needed to eight byte align a scatter/gather list CDB continuation descriptor. If the PAD LENGTH field is not set to zero in a scatter/gather list CDB continuation descriptor, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The CDB CONTINUATION DESCRIPTOR LENGTH field specifies the number of bytes that follow in this descriptor. The contents of the CDB CONTINUATION DESCRIPTOR LENGTH field shall be validated as described in 5.4.1.

Each scatter/gather list entry (see table 73) specifies the starting byte offset in the user object and number of bytes to be transferred to or from the command data buffer segment.

The first byte in the command data buffer segment is transferred to or from the user object byte offset indicated by the first scatter/gather list entry. Additional bytes are transferred to or from the command data buffer segment byte by byte until the number of bytes indicated by the first scatter/gather list entry have been transferred.

The next byte in the command data buffer segment (i.e., the first byte in the command data buffer segment that was not transferred by the first scatter/gather list entry) is transferred to the user object byte offset indicated by the second scatter/gather list entry. Additional bytes are transferred to or from the command data buffer segment byte by byte until the number of bytes indicated by the second scatter/gather list entry have been transferred.

This process is repeated until all the bytes indicated by the CDB LENGTH field have been transferred or all the scatter/gather list entries have been processed, which ever occurs first.

The description of scatter/gather list entry processing in this subclause describes the result achieved when GOOD status is returned. The device server may process the scatter/gather list entries in any order and transfer the results of that processing in any order that is consistent with the data buffer transfer requirements of the service delivery subsystem.

Each scatter/gather list entry has the format shown in table 73.

Table 73 — Scatter/gather list entry format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
7	USER OBJECT BYTE OFFSET _____ (LSB)							
8	(MSB) _____							
15	BYTES TO TRANSFER _____ (LSB)							

The USER OBJECT BYTE OFFSET field specifies the starting byte offset in the user object for this scatter/gather list entry.

The BYTES TO TRANSFER field specifies the number of bytes to transfer for this scatter/gather list entry.

It shall not be an error for the byte ranges specified by individual scatter/gather list entries to overlap.

If the values in the BYTES TO TRANSFER field and USER OBJECT BYTE OFFSET field result an attempt to read a byte that is beyond the user object logical length attribute value in the User Object Information attributes page (see 7.1.3.11), then:

- a) All bytes transferred by scatter/gather list entries prior to the scatter/gather list entry in which the error was detected shall be transferred;
- b) The bytes between the user object byte offset and the user object logical length shall be transferred;
- c) The command shall be terminated with CHECK CONDITION status, with the sense key shall be set to RECOVERED ERROR and the additional sense code set to READ PAST END OF USER OBJECT;
- d) The command-specific information sense data descriptor (see SPC-3) shall be included in the sense data; and
- e) The COMMAND-SPECIFIC INFORMATION field shall contain the number of bytes transferred by the command, including but not limited to the bytes transferred by this scatter/gather list entry.

5.4.3 Query list

The query list CDB continuation descriptor (see table 78) specifies the criteria for selecting the user objects whose User_Object_IDs are returned as matches by a QUERY command (see 6.26).

Table 74 — Query list CDB continuation descriptor format

Bit Byte	7	6	5	4	3	2	1	0
	CDB continuation descriptor header							
0	(MSB)	CDB CONTINUATION DESCRIPTOR TYPE (0002h)						(LSB)
1								
2	Reserved							
3	Reserved				PAD LENGTH (p-n)			
4	(MSB)	CDB CONTINUATION DESCRIPTOR LENGTH (n-7)						(LSB)
7								
	CDB continuation descriptor type specific data							
8	Reserved				QUERY TYPE			
9								
11	Reserved							
12	Query criteria entry (see table 76) [first]							
	⋮							
n	Query criteria entry (see table 76) [last]							
	CDB continuation descriptor alignment bytes							
n+1	Pad bytes (for eight-byte alignment)							
p								

The CDB CONTINUATION DESCRIPTOR TYPE field contains 0002h (i.e., query list CDB continuation descriptor).

The PAD LENGTH field specifies the number of bytes containing zeros that follow the CDB continuation descriptor type specific data.

The CDB CONTINUATION DESCRIPTOR LENGTH field contains the number of bytes of CDB continuation descriptor type specific data that follow in this descriptor. The contents of the CDB CONTINUATION DESCRIPTOR LENGTH field shall be validated as described in 5.4.1.

If the sum of the pad length and the CDB continuation descriptor length is not a multiple of eight, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The QUERY TYPE field (see table 75) specifies the format of the query criteria entries that follow.

Table 75 — QUERY TYPE field values

Code	Description
0h	A match with any query criteria entry shall cause the user object to appear in the list.
1h	Matching all query criteria entries shall cause the user object to appear in the list.
2h to Fh	Reserved

Each query criteria entry (see table 76) specifies matching criteria for one attribute.

Table 76 — Query criteria entry format

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
1								
2	(MSB)	QUERY ENTRY LENGTH (n-3)						
3								(LSB)
4	(MSB)	ATTRIBUTES PAGE						
7								(LSB)
8	(MSB)	ATTRIBUTE NUMBER						
11								(LSB)
12	(MSB)	MINIMUM ATTRIBUTE VALUE LENGTH (m-13)						
13								(LSB)
14	(MSB)	MINIMUM ATTRIBUTE VALUE						
m								(LSB)
m+1	(MSB)	MAXIMUM ATTRIBUTE VALUE LENGTH (n-m-2)						
m+2								(LSB)
m+3	(MSB)	MAXIMUM ATTRIBUTE VALUE						
n								(LSB)

The QUERY ENTRY LENGTH field specifies the number of bytes that follow in the query entry.

The ATTRIBUTES PAGE field specifies the page number of the attribute value. If the attributes page is not between 0h and 2FFF FFFFh, inclusive, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The ATTRIBUTE NUMBER field specifies the attribute number within the attributes page specified by the ATTRIBUTES PAGE field of the attribute value.

The MINIMUM ATTRIBUTE VALUE LENGTH field specifies the number of bytes that follow in the MINIMUM ATTRIBUTE VALUE field.

The MINIMUM ATTRIBUTE VALUE field specifies the minimum attribute value necessary for a user object to meet the criteria.

The MAXIMUM ATTRIBUTE VALUE LENGTH field specifies the number of bytes that follow in the MAXIMUM ATTRIBUTE VALUE field.

The MAXIMUM ATTRIBUTE VALUE field specifies the maximum attribute value necessary for a user object to meet the criteria.

5.4.4 User object

The user object CDB continuation descriptor (see table 77) specifies a user object to be processed by a command (e.g., a user object input to the READ MAPS AND COMPARE command (see 6.29)).

Table 77 — User object CDB continuation descriptor format

Bit Byte	7	6	5	4	3	2	1	0
	CDB continuation descriptor header							
0	(MSB)	CDB CONTINUATION DESCRIPTOR TYPE (0100h)						(LSB)
1								
2	Reserved							
3	Reserved				PAD LENGTH (000b)			
4	(MSB)	CDB CONTINUATION DESCRIPTOR LENGTH (n-7)						(LSB)
7								
	CDB continuation descriptor type specific data							
8	(MSB)	PARTITION_ID						(LSB)
15								
16	(MSB)	USER_OBJECT_ID						(LSB)
23								

The CDB CONTINUATION DESCRIPTOR TYPE field contains 0100h (i.e., user object CDB continuation descriptor).

The PAD LENGTH field is set to zero to indicate that no pad bytes are needed to eight byte align a user object CDB continuation descriptor. If the PAD LENGTH field is not set to zero in a user object CDB continuation descriptor, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The CDB CONTINUATION DESCRIPTOR LENGTH field contains the number of bytes that follow in this descriptor. The contents of the CDB CONTINUATION DESCRIPTOR LENGTH field shall be validated as described in 5.4.1.

The PARTITION_ID field specifies the Partition_ID (see 4.6.4) of the partition that contains the user object. If the partition identified by the PARTITION_ID field does not exist, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The USER_OBJECT_ID field specifies the User_Object_ID of the user object (see 4.6.5). If the user object identified by the USER_OBJECT_ID field does not exist, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

5.4.5 Copy user object source

The copy user object source CDB continuation descriptor (see table 78) specifies a user object as the source of bytes for a command (e.g., the source of bytes for a COPY USER OBJECTS command (see 6.h)).

Table 78 — Copy user object source CDB continuation descriptor format

Bit Byte	7	6	5	4	3	2	1	0
	CDB continuation descriptor header							
0	(MSB)	CDB CONTINUATION DESCRIPTOR TYPE (0101h)						(LSB)
1								
2	Reserved							
3	Reserved				PAD LENGTH (000b)			
4	(MSB)	CDB CONTINUATION DESCRIPTOR LENGTH (n-7)						(LSB)
7								
	CDB continuation descriptor type specific data							
8	(MSB)	SOURCE PARTITION_ID						(LSB)
15								
16	(MSB)	SOURCE USER_OBJECT_ID						(LSB)
23								
24	Reserved							CPY_ATTR
25	FREEZE	Reserved			TIME OF DUPLICATION			
26								
27	Reserved							

The CDB CONTINUATION DESCRIPTOR TYPE field contains 0101h (i.e., copy user object source CDB continuation descriptor).

The PAD LENGTH field is set to zero to indicate that no pad bytes are needed to eight byte align a copy user object source CDB continuation descriptor. If the PAD LENGTH field is not set to zero in a copy user object source CDB continuation descriptor, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The CDB CONTINUATION DESCRIPTOR LENGTH field contains the number of bytes that follow in this descriptor. The contents of the CDB CONTINUATION DESCRIPTOR LENGTH field shall be validated as described in 5.4.1.

The SOURCE PARTITION_ID field specifies the Partition_ID (see 4.6.4) of the partition that contains the source user object. If the partition identified by the SOURCE PARTITION_ID field does not exist, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The SOURCE USER_OBJECT_ID field specifies the User_Object_ID (see 4.6.5) of the source user object. If the user object identified by the SOURCE USER_OBJECT_ID field does not exist, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The CPY_ATTR (copy attributes) bit specifies whether the attributes from this source user object are copied to the destination user object. If the CPY_ATTR bit is set to zero, no attributes are copied from this source user object to the destination user object. If the CPY_ATTR bit is set to one, all application client settable attributes (see 7.1.3) are copied from this source user object to the destination user object.

If the CPY_ATTR bit is set to one and the reserved data space attribute in the User Object Information attributes page (see 7.1.3.11) of the source user object is set to a value other than zero, the reserved data space attribute in the destination user object shall be increased by the amount of data space reserved for the source user object (i.e., the copy operation shall be treated as appending the source user object to data already in the destination user object).

If the FREEZE bit is set to zero, the copy operation should not modify the contents of the object accessibility attribute in the User Object Information attributes page (see 7.1.3.11) of the source user object. If the FREEZE bit is set to one and source object freeze duplication management is supported (see 4.13.4.3), then the device server shall:

- 1) Set the object accessibility attribute in the User Object Information attributes page of the source user object to 0000 0001h before starting any copy operations that access the source user object; and
- 2) Restore the object accessibility attribute in the User Object Information attributes page of the source user object to its previous value after all copy operations that access the source user object are completed.

If the FREEZE bit is set to one and source object freeze duplication management is not supported, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The TIME OF DUPLICATION field specifies which time of duplication source object management method (see 4.13.4.2) applies to the source user object. If the TIME OF DUPLICATION field is set to DEFAULT (see table 44 in 4.13.4.2), then the default copy user objects time of duplication method attribute in the Partition Information attributes page (see 7.1.3.9) specifies which time of duplication source object management method applies to the source user object.

5.4.6 Extension capabilities CDB continuation descriptors

The extension capabilities CDB continuation descriptor (see table 79) adds one or more capabilities to the set of capabilities associated with the command.

Table 79 — Extension capabilities CDB continuation descriptor format

Bit Byte	7	6	5	4	3	2	1	0
	CDB continuation descriptor header							
0	(MSB)	CDB CONTINUATION DESCRIPTOR TYPE (FFEEh)						(LSB)
1								
2		Reserved						
3		Reserved				PAD LENGTH (000b)		
4	(MSB)	CDB CONTINUATION DESCRIPTOR LENGTH (k-7)						(LSB)
7								
	CDB continuation descriptor type specific data							
8		Extension capability (see 4.11.2.2)						
111		[first]						
		⋮						
k-103		Extension capability (see 4.11.2.2)						
k		[last]						

The CDB CONTINUATION DESCRIPTOR TYPE field contains FFEEh (i.e., extension capabilities CDB continuation descriptor).

The PAD LENGTH field is set to zero to indicate that no pad bytes are needed to eight byte align an extension capabilities CDB continuation descriptor. If the PAD LENGTH field is not set to zero in an extension capabilities CDB continuation descriptor, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The CDB CONTINUATION DESCRIPTOR LENGTH field contains the number of bytes that follow in this descriptor. The contents of the CDB CONTINUATION DESCRIPTOR LENGTH field shall be validated as described in 5.4.1.

Each extension capability is a capability (see 4.11.2.2) that adds to the object access capabilities of the command.

Unless the SECURITY METHOD field in the CDB specifies the NOSEC security method (i.e., if the CDB SECURITY METHOD field contains zero), the contents of the extension capabilities CDB continuation descriptor should be copied from a credential (see 4.12.5). The device server shall validate this credential as described in 4.12.6.1.

If a CDB continuation segment (see 5.3) contains more than one extension capabilities CDB continuation descriptor, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

6 Commands for OSD type devices

6.1 Summary of commands for OSD type devices

The commands for OSD type devices are listed in table 80. For the commands defined by this standard, the SERVICE ACTION field in the CDB uniquely identifies each command function. The referenced subclauses describe the service provided by each command function and the information that shall be passed to the OSD logical unit in order for it to perform that function.

Table 80 — Commands for OSD type devices (part 1 of 3)

Command name	Operation code	Service action ^a	Type	Reference
APPEND	7Fh	8887h	M	6.2
CLEAR	7Fh	8889h	M	6.3
COPY USER OBJECTS	7Fh	8893h	M	6.4
CREATE	7Fh	8882h	M	6.5
CREATE AND WRITE	7Fh	8892h	M	6.6
CREATE COLLECTION	7Fh	8895h	O ^b	6.8
CREATE CLONE	7Fh	88A8h	O ^c	6.7
CREATE PARTITION	7Fh	888Bh	M	6.9
CREATE SNAPSHOT	7Fh	88A9h	O ^d	6.10
CREATE USER TRACKING COLLECTION	7Fh	8894h	O ^b	6.11
DETACH CLONE	7Fh	88AAh	O ^c	6.12
FLUSH	7Fh	8888h	M	6.13
FLUSH COLLECTION	7Fh	889Ah	O ^b	6.14
FLUSH OSD	7Fh	889Ch	M	6.15
FLUSH PARTITION	7Fh	889Bh	M	6.16
FORMAT OSD	7Fh	8881h	O	6.17
GET ATTRIBUTES	7Fh	888Eh	M	6.18
Type Key: M = Command implementation is mandatory. O = Command implementation is optional.				
^a No entry in the service action column means that the SERVICE ACTION field does not apply to the command. Service action codes values between 8800h and 8F7Fh that are not listed in this table are reserved for future standardization. Service action code values between 8F80h and 8FFFh may have vendor specific command assignments. ^b Support for this command is mandatory if collections are supported (see 4.6.6). ^c Support for this command is mandatory if the maximum clones count attribute in the Root Information attributes page (see 7.1.3.8) is defined and contains a value other than zero. ^d Support for this command is mandatory if the maximum snapshots count attribute in the Root Information attributes page (see 7.1.3.8) is defined and contains a value other than zero. ^e Unless the security method in effect for the root object and every partition in the OSD logical unit is NOSEC (see 4.12.1), this command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID COMMAND OPERATION CODE. If the security method in effect for the root object or any partition in the OSD logical unit is not NOSEC, this command may be performed only by using the PERFORM SCSI COMMAND command (see 6.23). ^f The effects on established persistent reservations and registrations if the security method in effect for the root object or any partition changes from NOSEC to any other security method are described in 4.17.				

Table 80 — Commands for OSD type devices (part 2 of 3)

Command name	Operation code	Service action ^a	Type	Reference
GET MEMBER ATTRIBUTES	7Fh	88A2h	O ^b	6.19
INQUIRY	12h		M	SPC-4
LIST	7Fh	8883h	M	6.20
LIST COLLECTION	7Fh	8897h	O ^b	6.21
LOG SELECT ^e	4Ch		O	SPC-4
LOG SENSE ^e	4Dh		O	SPC-4
MODE SELECT(10) ^e	55h		O	SPC-4
MODE SENSE(10) ^e	5Ah		O	SPC-4
OBJECT STRUCTURE CHECK	7Fh	8880h	M	6.22
PERFORM SCSI COMMAND	7Fh	8F7Ch	M	6.23
PERFORM TASK MANAGEMENT FUNCTION	7Fh	8F7Dh	M	6.24
PERSISTENT RESERVE IN ^{e, f}	5Eh		O	SPC-4
PERSISTENT RESERVE OUT ^{e, f}	5Fh		O	SPC-4
PUNCH	7Fh	8884h	M	6.25
QUERY	7Fh	88A0h	O ^b	6.26
READ	7Fh	8885h	M	6.27
READ BUFFER ^e	3Ch		O	SPC-4
READ MAP	7Fh	88B1h	M	6.28
READ MAPS AND COMPARE	7Fh	88B2h	M	6.29
RECEIVE DIAGNOSTIC RESULTS ^e	1Ch		O	SPC-4
REFRESH SNAPSHOT OR CLONE	7Fh	88ABh	O ^{c, d}	6.30
REMOVE	7Fh	888Ah	M	6.31
REMOVE COLLECTION	7Fh	8896h	O	6.32
REMOVE MEMBER OBJECTS	7Fh	88A1h	O ^b	6.33
REMOVE PARTITION	7Fh	888Ch	M	6.34
Type Key: M = Command implementation is mandatory. O = Command implementation is optional.				
^a No entry in the service action column means that the SERVICE ACTION field does not apply to the command. Service action codes values between 8800h and 8F7Fh that are not listed in this table are reserved for future standardization. Service action code values between 8F80h and 8FFFh may have vendor specific command assignments. ^b Support for this command is mandatory if collections are supported (see 4.6.6). ^c Support for this command is mandatory if the maximum clones count attribute in the Root Information attributes page (see 7.1.3.8) is defined and contains a value other than zero. ^d Support for this command is mandatory if the maximum snapshots count attribute in the Root Information attributes page (see 7.1.3.8) is defined and contains a value other than zero. ^e Unless the security method in effect for the root object and every partition in the OSD logical unit is NOSEC (see 4.12.1), this command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID COMMAND OPERATION CODE. If the security method in effect for the root object or any partition in the OSD logical unit is not NOSEC, this command may be performed only by using the PERFORM SCSI COMMAND command (see 6.23). ^f The effects on established persistent reservations and registrations if the security method in effect for the root object or any partition changes from NOSEC to any other security method are described in 4.17.				

Table 80 — Commands for OSD type devices (part 3 of 3)

Command name	Operation code	Service action ^a	Type	Reference
REPORT LUNS	A0h		M	SPC-4
REPORT SUPPORTED OPERATION CODES ^e	A3h	0Ch	O	SPC-4
REPORT SUPPORTED TASK MANAGEMENT FUNCTIONS ^e	A3h	0Dh	O	SPC-4
REPORT TARGET PORT GROUPS ^e	A3h	0Ah	O	SPC-4
REQUEST SENSE	03h		M	SPC-4
RESTORE PARTITION FROM SNAPSHOT	7Fh	88ACh	O ^d	6.35
SEND DIAGNOSTIC ^e	1Dh		M	SPC-4
SET ATTRIBUTES	7Fh	888Fh	M	6.36
SET KEY	7Fh	8898h	M	6.37
SET MASTER KEY	7Fh	8899h	M	6.38
SET MEMBER ATTRIBUTES	7Fh	88A3h	O ^b	6.39
SET TARGET PORT GROUPS ^e	A4h	0Ah	O	SPC-4
TEST UNIT READY	00h		M	SPC-4
WRITE	7Fh	8886h	M	6.40
WRITE BUFFER ^e	3Bh		O	SPC-4
Type Key: M = Command implementation is mandatory. O = Command implementation is optional.				
^a No entry in the service action column means that the SERVICE ACTION field does not apply to the command. Service action codes values between 8800h and 8F7Fh that are not listed in this table are reserved for future standardization. Service action code values between 8F80h and 8FFFh may have vendor specific command assignments. ^b Support for this command is mandatory if collections are supported (see 4.6.6). ^c Support for this command is mandatory if the maximum clones count attribute in the Root Information attributes page (see 7.1.3.8) is defined and contains a value other than zero. ^d Support for this command is mandatory if the maximum snapshots count attribute in the Root Information attributes page (see 7.1.3.8) is defined and contains a value other than zero. ^e Unless the security method in effect for the root object and every partition in the OSD logical unit is NOSEC (see 4.12.1), this command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID COMMAND OPERATION CODE. If the security method in effect for the root object or any partition in the OSD logical unit is not NOSEC, this command may be performed only by using the PERFORM SCSI COMMAND command (see 6.23). ^f The effects on established persistent reservations and registrations if the security method in effect for the root object or any partition changes from NOSEC to any other security method are described in 4.17.				

6.2 APPEND

The APPEND command (see table 81) causes the specified number of bytes to be written to the designated object starting immediately after the user object's logical length.

Table 81 — APPEND command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (8887h)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID							
24	(MSB)							
31	USER_OBJECT_ID							
32	(MSB)							
39	LENGTH							
40	Reserved							
47								
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

The contents of the LENGTH field are described in 5.2.9. The data to be written to the user object shall be placed in the Data-Out Buffer as described in 4.15.4.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

If an APPEND command causes the value in the user object logical length attribute in the User Object Information attributes page (see 7.1.3.11) to exceed the value in the maximum user object length attribute in the User Object Quotas attributes page (see 7.1.3.14), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the maximum user object length quota.

If an APPEND command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the capacity quota.

6.3 CLEAR

The CLEAR command (see table 82) causes the specified number of bytes containing zero to be written to the specified user object at the specified relative location.

Table 82 — CLEAR command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (8889h)							
10	(LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13								
15	Reserved							
16	(MSB)							
23	PARTITION_ID							
24	(LSB)							
24	(MSB)							
31	USER_OBJECT_ID							
32	(LSB)							
32	(MSB)							
39	CLEAR LENGTH							
40	(LSB)							
40	(MSB)							
47	CLEAR STARTING BYTE ADDRESS							
48	(LSB)							
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
52	(LSB)							
52								
79	Get and set attributes parameters (see 5.2.6)							
80								
80								
183	Capability (see 5.2.4)							
184								
184								
235	Security parameters (see 5.2.11)							
235								

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

The CLEAR LENGTH field specifies the number of bytes containing zero to be written to the specified user object.

The CLEAR STARTING BYTE ADDRESS field specifies the location where the writing of bytes containing zero is to commence relative to the first byte (i.e., byte zero) of the specified user object.

Writing zero to a byte at a location that is greater than or equal to the value in the user object logical length attribute in the User Object Information attributes page (see 7.1.3.11) shall implicitly increase the value in the user object logical length attribute to the largest location of any byte written.

The command data segment of the Data-Out Buffer is not used by the CLEAR command.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

If a CLEAR command causes the value in the user object logical length attribute in the User Object Information attributes page (see 7.1.3.11) to exceed the value in the maximum user object length attribute in the User Object Quotas attributes page (see 7.1.3.14), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the maximum user object length quota.

If a CLEAR command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated. The quota testing principles described in 4.10.3 apply to the testing of the maximum capacity quota.

6.4 COPY USER OBJECTS

The COPY USER OBJECTS command (see table 85) causes the OSD device server to allocate and initialize one user object and then copy data from one source user objects to the newly created user object.

Table 83 — COPY USER OBJECTS command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8893h) _____ (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
14	DUPLICATION METHOD							
15	Reserved							
16	(MSB) _____							
23	DESTINATION PARTITION_ID _____ (LSB)							
24	(MSB) _____							
31	REQUESTED DESTINATION USER_OBJECT_ID _____ (LSB)							
32	Reserved							
47	Reserved							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	Reserved							
79	Get and set attributes parameters (see 5.2.6) _____							
80	Reserved							
183	Capability (see 5.2.4) _____							
184	Reserved							
235	Security parameters (see 5.2.11) _____							

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The DUPLICATION METHOD field specifies which duplication method (see 4.13.3) applies to the COPY USER OBJECTS command. If the DUPLICATION METHOD field is set to DEFAULT (see table 43 in 4.13.3), then the default copy user objects duplication method attribute in the Partition Information attributes page (see 7.1.3.9) specifies which duplication method applies to the COPY USER OBJECTS command.

The contents of the DESTINATION PARTITION_ID field are described in 5.2.10. If the DESTINATION PARTITION_ID field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the REQUESTED DESTINATION USER_OBJECT_ID field specify the User_Object_ID (see 4.6.5) to be assigned to the created user object that is to serve as the destination for the data copies from the source user object. If the REQUESTED DESTINATION USER_OBJECT_ID field contains zero, any User_Object_ID may be assigned. If the REQUESTED DESTINATION USER_OBJECT_ID field contains any value other than zero and the device server is unable to assign the requested User_Object_ID to the created user object, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

Within a partition, the device server shall not allow:

- a) The same User_Object_ID to be associated with more than one user object at any point in time; or
- b) A User_Object_ID to have the same value as any assigned Collection_Object_ID.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if the CDB continuation segment (see 5.3):

- a) Does not contain any copy user object source CDB continuation descriptors (see 5.4.5);
- b) Contains more than one copy user object source CDB continuation descriptor;
- c) Contains more than one extension capabilities CDB continuation descriptor (see 5.4.6); or
- d) Contains any CDB continuation descriptors other than the following:
 - A) Copy user object source CDB continuation descriptor (see 5.4.5); and
 - B) Extension capabilities CDB continuation descriptor (see 5.4.6).

The copy user object source CDB continuation descriptor specifies the data that is to be written to the created user object.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.14. The User_Object_ID assigned by the COPY USER OBJECTS command may be obtained from the Current Command attributes page (see 7.1.3.31).

The capability is described in 5.2.4. The COPY USER OBJECTS command accesses two user objects. One capability is necessary for each user object accessed. One capability appears in the CDB. The other capability appears in the CDB continuation segment (see 5.3).

The security parameters are described in 5.2.11.

The assigned User_Object_ID shall be placed in the Collection_Object_ID or User_Object_ID attribute in the Current Command attributes page (see 7.1.3.31).

If a COPY USER OBJECTS command causes the value in the number of collections and user objects attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the object count attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the object count quota.

If a COPY USER OBJECTS command causes the value in the user object logical length attribute in the User Object Information attributes page (see 7.1.3.11) to exceed the value in the maximum user object length attribute in the User Object Quotas attributes page (see 7.1.3.14), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the maximum user object length quota.

If a COPY USER OBJECTS command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the capacity quota.

6.5 CREATE

The CREATE command (see table 84) causes the OSD device server to allocate and initialize one or more user objects.

Table 84 — CREATE command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8882h) _____ (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	(MSB) _____							
31	REQUESTED USER_OBJECT_ID _____ (LSB)							
32	(MSB) _____							
33	NUMBER OF USER OBJECTS _____ (LSB)							
34	_____							
47	Reserved _____							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10. If the PARTITION_ID field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the REQUESTED USER_OBJECT_ID field specify the User_Object_ID (see 4.6.5) to be assigned to the created user object. If the REQUESTED USER_OBJECT_ID field contains zero, any User_Object_ID may be assigned. If the REQUESTED USER_OBJECT_ID field contains any value other than zero and the device server is unable to assign the requested User_Object_ID to the created user object, the user object shall not be created and the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

Within a partition, the device server shall not allow:

- a) The same User_Object_ID to be associated with more than one user object at any point in time; or
- b) A User_Object_ID to have the same value as any assigned Collection_Object_ID.

The NUMBER OF USER OBJECTS field specifies the number of user objects to be created. If the NUMBER OF USER OBJECTS field contains zero or one, one user object shall be created. Otherwise:

- a) The number of user objects created shall equal the value in the NUMBER OF USER OBJECTS field;
- b) The user objects created shall be assigned consecutive valued User_Object_IDs; and
- c) The lowest valued User_Object_ID shall be placed in the Collection_Object_ID or User_Object_ID attribute of the Current Command attributes page (see 7.1.3.31).

If the NUMBER OF USER OBJECTS field contains a value that is greater than one and the requested User_Object_ID is not zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB if:

- a) If the NUMBER OF USER OBJECT field contains a value that is greater than one;
- b) The GET/SET CDBFMT field contains 10b; and
- c) The GET ATTRIBUTES PAGE field (see 5.2.6.3) contains a value other than FFFF FFFEh (i.e., the Current Command attributes page).

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the get and set attributes parameters request the retrieval of attributes from pages other than the Current Command attributes page, the attributes for every created user object shall be returned using list type Fh (see 7.1.4).

| The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

| The lowest valued assigned User_Object_ID shall be placed in the Collection_Object_ID or User_Object_ID attribute in the Current Command attributes page (see 7.1.3.31).

If a CREATE command causes the value in the number of collections and user objects attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the object count attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the object count quota.

If a CREATE command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the capacity quota.

6.6 CREATE AND WRITE

The CREATE AND WRITE command (see table 85) causes the OSD device server to allocate and initialize one user object and then write data to the newly created user object.

Table 85 — CREATE AND WRITE command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (8892h)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID							
24	(MSB)							
31	REQUESTED USER_OBJECT_ID							
32	(MSB)							
39	LENGTH							
40	(MSB)							
47	STARTING BYTE ADDRESS							
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10. If the PARTITION_ID field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the REQUESTED_USER_OBJECT_ID field specify the User_Object_ID (see 4.6.5) to be assigned to the created user object. If the REQUESTED_USER_OBJECT_ID field contains zero, any User_Object_ID may be assigned. If the REQUESTED_USER_OBJECT_ID field contains any value other than zero and the device server is unable to assign the requested User_Object_ID to the created user object, the user object shall not be created and the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

Within a partition, the device server shall not allow:

- a) The same User_Object_ID to be associated with more than one user object at any point in time; or
- b) A User_Object_ID to have the same value as any assigned Collection_Object_ID.

The contents of the LENGTH field are described in 5.2.9. The data to be written to the user object shall be placed in the Data-Out Buffer as described in 5.2.9.

The contents of the STARTING_BYTE_ADDRESS field are described in 5.2.12. If the CDB continuation segment (see 5.3), if any, contains a scatter/gather list CDB continuation descriptor and the STARTING_BYTE_ADDRESS field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the CDB_CONTINUATION_LENGTH field are described in 5.2.5. If the CDB_CONTINUATION_LENGTH field is not set to zero and the CDB continuation segment (see 5.3) contains a scatter/gather list CDB continuation descriptor, that descriptor shall be processed as described in 5.4.2.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if the CDB continuation segment (see 5.3):

- a) Contains more than one scatter/gather list CDB continuation descriptor; or
- b) Contains any CDB continuation descriptors other than the scatter/gather list CDB continuation descriptor.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15. The User_Object_ID assigned by the CREATE AND WRITE command may be obtained from the Current Command attributes page (see 7.1.3.31).

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

The assigned User_Object_ID shall be placed in the Collection_Object_ID or User_Object_ID attribute in the Current Command attributes page (see 7.1.3.31).

If a CREATE AND WRITE command causes the value in the number of collections and user objects attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the object count attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the object count quota.

If a CREATE AND WRITE command causes the value in the user object logical length attribute in the User Object Information attributes page (see 7.1.3.11) to exceed the value in the maximum user object length attribute in the User Object Quotas attributes page (see 7.1.3.14), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the maximum user object length quota.

If a CREATE AND WRITE command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes

page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the capacity quota.

6.7 CREATE CLONE

6.7.1 Introduction

The CREATE CLONE command (see table 86) causes the OSD device server to allocate and initialize a destination partition as a clone partition (see 4.13.2) and then copy all user objects, collections, and attributes from a source partition to the newly created clone partition.

Table 86 — CREATE CLONE command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (88A8h) _____ (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	IMMED_TR	Reserved	GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	FREEZE	Reserved			TIME OF DUPLICATION			
14	DUPLICATION METHOD							
15	Reserved							
16	(MSB) _____							
23	SOURCE PARTITION_ID _____ (LSB)							
24	(MSB) _____							
31	REQUESTED DESTINATION PARTITION_ID _____ (LSB)							
32	Reserved _____							
47								
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	Get and set attributes parameters (see 5.2.6) _____							
79								
80	Capability (see 5.2.4) _____							
183								
184	Security parameters (see 5.2.11) _____							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The IMMED_TR bit is described in 5.2.5.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

If the FREEZE bit is set to zero, the CREATE CLONE command shall not modify the contents of the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) of the source partition. If the FREEZE bit is set to one and source object freeze duplication management is supported (see 4.13.4.3), then the device server shall modify the contents of the object accessibility attribute in the Partition Information attributes page of the source partition as described in 6.7.2 and 6.7.4.

The TIME OF DUPLICATION field specifies which time of duplication source object management method (see 4.13.4.2) applies to the CREATE CLONE command. If the TIME OF DUPLICATION field is set to DEFAULT (see table 44 in 4.13.4.2), then the default clone time of duplication method attribute in the Partition Information attributes page (see 7.1.3.9) of the source partition specifies which time of duplication management method applies to the CREATE CLONE command.

The DUPLICATION METHOD field specifies which duplication method (see 4.13.3) applies to the CREATE CLONE command. If the DUPLICATION METHOD field is set to DEFAULT (see table 43 in 4.13.3), then the default clone duplication method attribute in the Partition Information attributes page (see 7.1.3.9) of the source partition specifies which duplication method applies to the CREATE CLONE command.

The SOURCE PARTITION_ID field contains the Partition_ID (see 4.6.4) of the source partition for the CREATE CLONE command.

The contents of the REQUESTED DESTINATION PARTITION_ID field specify the Partition_ID to be assigned to the created clone partition. If the REQUESTED DESTINATION PARTITION_ID field contains zero any Partition_ID may be assigned. If the REQUESTED DESTINATION PARTITION_ID field contains any value other than zero and the device server is unable to assign the requested Partition_ID to the destination partition, the destination partition shall not be created and the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if the CDB continuation segment (see 5.3):

- a) Does not contain one extension capabilities CDB continuation descriptor (see 5.4.6); or
- b) Contains any CDB continuation descriptors other than the extension capabilities CDB continuation descriptor.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15. The destination Partition_ID assigned by the CREATE SNAPSHOT command may be obtained from the Current Command attributes page (see 7.1.3.31).

The capability is described in 5.2.4. The CREATE CLONE command accesses two partitions. One capability is necessary for each partition accessed. One capability appears in the CDB. The other capability appears in the CDB continuation segment (see 5.3).

The security parameters are described in 5.2.11.

The CREATE CLONE command does not initialize the partition key or the working keys (see 4.12.9.1) for the destination partition. Proper operation of any security method other than NOSEC (see 4.12.4) requires that the following commands be processed without errors before other commands are addressed to the destination partition:

- a) A SET KEY command (see 6.37) that establishes the partition key; and
- b) One or more SET KEY commands that establish one or more working keys for the partition.

A CREATE CLONE command whose capability (see 4.11.2.2) for the destination partition has the SET_ATTR bit set to one and POL/SEC bit set to one is allowed to avoid the need for SET KEY commands by setting the default security method attribute to NOSEC in the Partition Policy/Security attributes page (see 7.1.3.23) for the created partition.

If the REQUESTED DESTINATION PARTITION_ID field is not set to zero, SET KEY commands are not needed to enable the tracking of the progress of a CREATE CLONE command with the IMMED_TR bit set to one in the following cases:

- a) If the READ permission bit is set to one in the capability that allowed creation of the destination partition, that capability may be used in LIST COLLECTION commands (see 6.21) that list the contents of the snapshot/clone tracking well known collection (see 4.6.6.5.3); and
- b) If the GET_ATTR permission bit is set to one in the capability that allowed creation of the destination partition, that capability may be used in GET ATTRIBUTES commands (see 6.18) or equivalents that retrieve attributes from the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection.

6.7.2 Processing before the IMMED_TR bit takes effect

A CREATE CLONE command shall not be completed with GOOD status until at least all the operations described in this subclause have been performed. These operations shall be performed before completing the command with GOOD status even if the IMMED_TR bit is set to one.

If the FREEZE bit is set to one and source object freeze duplication management (see 4.13.4.3) is not supported, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD CDB.

If the requested time of duplication source object management method (see 4.13.4.2) is not supported or the requested duplication method (see 4.13.3) is not supported, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the SOURCE PARTITION_ID field contains zero or the Partition_ID (see 4.6.4) of a partition that does not exist, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB, if any of the follow conditions exist in the attribute values in the Snapshots Information attributes page (see 7.1.3.30) of the source partition:

- a) The partition type attribute contains 00h (i.e., primary partition);
- b) The partition type attribute contains 02h (i.e., clone partition);
- c) The clones count attribute contains a value that is equal to the value in the maximum clones count attribute in the Root Information attributes page (see 7.1.3.8); or
- d) The branch depth attribute contains a value that is equal to the value in the maximum branch depth attribute in the Root Information attributes page.

If the REQUESTED DESTINATION PARTITION_ID field contains any value other than zero and the device server is unable to assign the requested Partition_ID to the created partition, the partition shall not be created and the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The device server shall not allow the same Partition_ID to be associated with more than one partition at any point in time.

If a CREATE CLONE command causes the value in the number of partitions attribute in the Root Information attributes page (see 7.1.3.8) to exceed the value in the partition count attribute in the Root Quotas attributes page (see 7.1.3.12), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the partition count quota.

The device server shall create the requested destination partition and initialize it as if a CREATE PARTITION command (see 6.9) were being processed.

The assigned Partition_ID shall be placed in the Partition_ID attribute in the Current Command attributes page (see 7.1.3.31). The Collection_Object_ID or User_Object_ID attribute in the Current Command attributes page shall be set to zero.

The object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the destination partition shall be set to 0000 0001h.

The snapshot/clone tracking well known collection (see 4.6.6.5.3) shall be created in the destination partition, and initialized, including at least the following:

- a) Every user object and collection in the source partition shall have their User_Object_ID (see 4.6.5) or Collection_Object_ID (see 4.6.6) inserted as a member of the TRACKING collection (see 4.6.6.3); and
- b) The Command Tracking attributes page (see 7.1.3.20) shall be initialized to include at least the following:
 - A) The percent complete attribute shall be set to zero;
 - B) The active command status attribute shall be set to 88A8h (i.e., CREATE CLONE command in progress); and
 - C) The ended command status attribute shall be set to FFFFh.

The following attributes in the Snapshots Information attributes page (see 7.1.3.30) of the source partition shall be set as follows:

- a) One of the clone destination attributes that is undefined (see 3.1.51) shall be defined and set to the Partition_ID (see 4.6.4) of the destination partition;
- b) If it is defined (see 3.1.14), the clones count attribute shall have its value incremented by one. If the clones count attribute is undefined (see 3.1.51), then it shall be defined and set to a value of one;
- c) The create completion time attribute shall be made undefined (see 3.1.51); and
- d) The refresh completion time attribute shall be made undefined.

The following attributes in the Snapshots Information attributes page (see 7.1.3.30) of the destination partition shall be set as follows:

- a) The partition type attribute shall be set to 02h (i.e., clone partition);
- b) The source partition attribute shall be set to the Partition_ID (see 4.6.4) of the source partition; and
- c) The branch depth attribute shall be set as follows:
 - A) If the branch depth attribute is defined (see 3.1.14) in the Snapshots Information attributes page of the source partition, then the branch depth attribute value for the destination partition shall be set to one plus the value in the branch depth attribute for the source partition; or

- B) If the branch depth attribute is undefined (see 3.1.51) in the Snapshots Information attributes page of the source partition, then the branch depth attribute value for the destination partition shall be set to one.

If the FREEZE bit is set to one, the device server shall:

- a) Note the value of the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the source partition for use in 6.7.4; and
- b) Set the object accessibility attribute in the Partition Information attributes page in the source partition to 0000 0001h.

6.7.3 Processing after the IMMED_TR bit takes effect, if any

Every user object and collection in the source partition shall be duplicated in the destination clone partition using the:

- a) Duplication method (see 4.13.3) specified by the CDB; and
- b) Time of duplication method (see 4.13.4.2) specified by the CDB.

The membership and attributes of the snapshot/clone tracking well known collection for the destination partition should be maintained to restarting of an interrupted CREATE CLONE command with the minimum of repeated work (e.g., user objects or collections that have been fully duplicated should be removed from the snapshot/clone tracking well known collection). Other factors (e.g., meeting the requirements of the END time of duplication method (see 4.13.4.2)) may cause user objects and collections to be added to the snapshot/clone tracking well known collection.

6.7.4 Command completion

When an error is encountered or when all user objects and collections in the source partition have been duplicated in the destination clone partition as described in 6.7.3, the CREATE CLONE command processing shall be completed as described in this subclause.

If the FREEZE bit is set to one, the device server shall restore the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the source partition to the value noted in 6.7.2.

At least the following changes shall be made in the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the destination partition:

- a) The active command status attribute shall be set to zero;
- b) The ended command status attribute shall be set to indicate the condition (e.g., success or error) of the CREATE CLONE command processing; and
- c) If sense data is available, it shall be placed in the sense data attribute.

If the CREATE CLONE command processing is complete (i.e., if the percent complete attribute in the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the destination partition is set to 100) and the ended command status attribute in the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the destination partition has been set to 0000h (i.e., GOOD status command completion), then:

- a) The create completion time attribute in the Snapshots Information attributes page (see 7.1.3.30) in the destination clone partition shall be set to the value of the clock attribute in the Root Information attributes page (see 7.1.3.8); and
- b) The object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the destination clone partition shall be set to 0000 0000h (i.e., allow all accesses).

If the IMMED_TR bit is set to zero, status shall be returned for the CREATE CLONE command.

6.8 CREATE COLLECTION

The CREATE COLLECTION command (see table 87) initializes a new LINKED collection (see 4.6.6.2).

Table 87 — CREATE COLLECTION command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (8895h)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID							
24	(MSB)							
31	REQUESTED COLLECTION_OBJECT_ID							
32	Reserved							
47								
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field (see 5.2.10) specify the Partition_ID of the partition in which the collection is to be created.

The contents of the REQUESTED COLLECTION_OBJECT_ID field specify the Collection_Object_ID (see 4.6.6) to be assigned to the created collection. If the REQUESTED COLLECTION_OBJECT_ID field contains zero any Collection_Object_ID may be assigned. If the REQUESTED COLLECTION_OBJECT_ID field contains any value other than zero and the device server is unable to assign the requested Collection_Object_ID to the created collection, the collection shall not be created and the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

Within a partition, the device server shall not allow:

- a) The same Collection_Object_ID to be associated with more than one collection at any point in time; or
- b) A Collection_Object_ID to have the same value as any assigned User_Object_ID.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15. The Collection_Object_ID assigned by the CREATE COLLECTION command may be obtained from the Current Command attributes page (see 7.1.3.31).

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

The collection type attribute in the Collection Information attributes page (see 7.1.3.10) shall be set to 00h (i.e., LINKED).

The assigned Collection_Object_ID shall be placed in the Collection_Object_ID or User_Object_ID attribute in the Current Command attributes page (see 7.1.3.31).

If a CREATE COLLECTION command causes the value in the number of collections and user objects attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the object count attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the object count quota.

If a CREATE COLLECTION command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the capacity quota.

6.9 CREATE PARTITION

The CREATE PARTITION command (see table 88) allocates and initializes a new partition.

Table 88 — CREATE PARTITION command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (888Bh) _____ (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved _____							
15	Reserved _____							
16	(MSB) _____							
23	REQUESTED PARTITION_ID _____ (LSB)							
24	Reserved _____							
47	Reserved _____							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	Get and set attributes parameters (see 5.2.6) _____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	Capability (see 5.2.4) _____							
183	Capability (see 5.2.4) _____							
184	Security parameters (see 5.2.11) _____							
235	Security parameters (see 5.2.11) _____							

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the REQUESTED PARTITION_ID field specify the Partition_ID (see 4.6.4) to be assigned to the created partition. If the REQUESTED PARTITION_ID field contains zero any Partition_ID may be assigned. If the REQUESTED PARTITION_ID field contains any value other than zero and the device server is unable to assign the requested Partition_ID to the created partition, the partition shall not be created and the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The device server shall not allow the same Partition_ID to be associated with more than one partition at any point in time.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15. The Partition_ID assigned by the CREATE PARTITION command may be obtained from the Current Command attributes page (see 7.1.3.31).

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

The assigned Partition_ID shall be placed in the Partition_ID attribute in the Current Command attributes page (see 7.1.3.31). The Collection_Object_ID or User_Object_ID attribute in the Current Command attributes page shall be set to zero.

If a CREATE PARTITION command causes the value in the number of partitions attribute in the Root Information attributes page (see 7.1.3.8) to exceed the value in the partition count attribute in the Root Quotas attributes page (see 7.1.3.12), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the partition count quota.

The CREATE PARTITION command does not initialize the partition key or the working keys (see 4.12.9.1) for the new partition. Proper operation of any security method other than NOSEC (see 4.12.4) requires that the following commands be processed without errors before other commands are addressed to the partition:

- a) A SET KEY command (see 6.37) that establishes the partition key; and
- b) One or more SET KEY commands that establish one or more working keys for the partition.

A CREATE PARTITION command whose capability (see 4.11.2.2) has the SET_ATTR bit set to one and POL/SEC bit set to one is allowed to avoid the need for SET KEY commands by setting the default security method attribute to NOSEC in the Partition Policy/Security attributes page (see 7.1.3.23) for the created partition.

6.10 CREATE SNAPSHOT

6.10.1 Introduction

The CREATE SNAPSHOT command (see table 89) causes the OSD device server to allocate and initialize a destination partition as a snapshot partition (see 4.13.2) and then copy all user objects, collections, and attributes from a source partition to the newly created snapshot partition.

Table 89 — CREATE SNAPSHOT command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (88A9h)							
10	Reserved			DPO	FUA	ISOLATION		
11	IMMED_TR	Reserved	GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	FREEZE	Reserved			TIME OF DUPLICATION			
14	DUPLICATION METHOD							
15	Reserved							
16	(MSB)							
23	SOURCE PARTITION_ID							
24	(MSB)							
31	REQUESTED DESTINATION PARTITION_ID							
32	Reserved							
47								
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The IMMEDIATE bit is described in 5.2.5.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

If the FREEZE bit is set to zero, the CREATE SNAPSHOT command shall not modify the contents of the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) of the source partition. If the FREEZE bit is set to one and source object freeze duplication management is supported (see 4.13.4.3), then the device server shall modify the contents of the object accessibility attribute in the Partition Information attributes page of the source partition as described in 6.10.2 and 6.10.4.

The TIME OF DUPLICATION field specifies which time of duplication source object management method (see 4.13.4.2) applies to the CREATE SNAPSHOT command. If the TIME OF DUPLICATION field is set to DEFAULT (see table 44 in 4.13.4.2), then the default snapshot time of duplication method attribute in the Partition Information attributes page (see 7.1.3.9) of the source partition specifies which time of duplication management method applies to the CREATE SNAPSHOT command.

The DUPLICATION METHOD field specifies which duplication method (see 4.13.3) applies to the CREATE SNAPSHOT command. If the DUPLICATION METHOD field is set to DEFAULT (see table 43 in 4.13.3), then the default snapshot duplication method attribute in the Partition Information attributes page (see 7.1.3.9) of the source partition specifies which duplication method applies to the CREATE SNAPSHOT command.

The SOURCE PARTITION_ID field contains the Partition_ID (see 4.6.4) of the source partition for the CREATE SNAPSHOT command.

The contents of the REQUESTED DESTINATION PARTITION_ID field specify the Partition_ID to be assigned to the created snapshot partition. If the REQUESTED DESTINATION PARTITION_ID field contains zero any Partition_ID may be assigned. If the REQUESTED DESTINATION PARTITION_ID field contains any value other than zero and the device server is unable to assign the requested Partition_ID to the destination partition, the destination partition shall not be created and the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if the CDB continuation segment (see 5.3):

- a) Does not contain one extension capabilities CDB continuation descriptor (see 5.4.6); or
- b) Contains any CDB continuation descriptors other than the extension capabilities CDB continuation descriptor.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15. The destination Partition_ID assigned by the CREATE SNAPSHOT command may be obtained from the Current Command attributes page (see 7.1.3.31).

The capability is described in 5.2.4. The CREATE SNAPSHOT command accesses two partitions. One capability is necessary for each partition accessed. One capability appears in the CDB. The other capability appears in the CDB continuation segment (see 5.3).

The security parameters are described in 5.2.11.

The CREATE SNAPSHOT command does not initialize the partition key or the working keys (see 4.12.9.1) for the destination partition. Proper operation of any security method other than NOSEC (see 4.12.4) requires that the following commands be processed without errors before other commands are addressed to the destination partition:

- a) A SET KEY command (see 6.37) that establishes the partition key; and
- b) One or more SET KEY commands that establish one or more working keys for the partition.

A CREATE SNAPSHOT command whose capability (see 4.11.2.2) for the destination partition has the SET_ATTR bit set to one and POL/SEC bit set to one is allowed to avoid the need for SET KEY commands by setting the default security method attribute to NOSEC in the Partition Policy/Security attributes page (see 7.1.3.23) for the created partition.

If the REQUESTED DESTINATION PARTITION_ID field is not set to zero, SET KEY command are not needed to enable the tracking of the progress of a CREATE SNAPSHOT command with the IMMED_TR bit set to one in the following cases:

- a) If the READ permission bit is set to one in the capability that allowed creation of the destination partition, that capability may be used in LIST COLLECTION commands (see 6.21) that list the contents of the snapshot/clone tracking well known collection (see 4.6.6.5.3); and
- b) If the GET_ATTR permission bit is set to one in the capability that allowed creation of the destination partition, that capability may be used in GET ATTRIBUTES commands (see 6.18) or equivalents that retrieve attributes from the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection.

6.10.2 Processing before the IMMED_TR bit takes effect

A CREATE SNAPSHOT command shall not be completed with GOOD status until at least all the operations described in this subclause have been performed. These operations shall be performed before completing the command with GOOD status even if the IMMED_TR bit is set to one.

If the FREEZE bit is set to one and source object freeze duplication management (see 4.13.4.3) is not supported, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the requested time of duplication source object management method (see 4.13.4.2) is not supported or the requested duplication method (see 4.13.3) is not supported, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the SOURCE PARTITION_ID field contains zero or the Partition_ID (see 4.6.4) of a partition that does not exist, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB, if any of the following conditions exist in the attribute values in the Snapshots Information attributes page (see 7.1.3.30) of the source partition:

- a) The partition type attribute contains 01h (i.e., snapshot partition); or
- b) The snapshots count attribute contains a value that is equal to the value in the maximum snapshots count attribute in the Root Information attributes page (see 7.1.3.8).

If the REQUESTED DESTINATION PARTITION_ID field contains any value other than zero and the device server is unable to assign the requested Partition_ID to the created partition, the partition shall not be created and the command

shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The device server shall not allow the same Partition_ID to be associated with more than one partition at any point in time.

If a CREATE SNAPSHOT command causes the value in the number of partitions attribute in the Root Information attributes page (see 7.1.3.8) to exceed the value in the partition count attribute in the Root Quotas attributes page (see 7.1.3.12), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the partition count quota.

The device server shall create the requested destination partition and initialize it as if a CREATE PARTITION command (see 6.9) were being processed.

The assigned Partition_ID shall be placed in the Partition_ID attribute in the Current Command attributes page (see 7.1.3.31). The Collection_Object_ID or User_Object_ID attribute in the Current Command attributes page shall be set to zero.

The object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the destination partition shall be set to 0000 0001h.

If it is defined (see 3.1.14), the snapshots count attribute in the Snapshots Information attributes page (see 7.1.3.30) of the source partition shall have its valued incremented by one. If the snapshots count attribute in the Snapshots Information attributes page of the source partition is undefined (see 3.1.51), then it shall be defined and set to a value of one.

The following attributes in the Snapshots Information attributes page (see 7.1.3.30) of the destination partition shall be set as follows:

- a) The partition type attribute shall be set to 01h (i.e., snapshot partition);
- b) The source partition attribute shall be set to the Partition_ID (see 4.6.4) of the source partition; and
- c) The branch depth attribute shall be set as follows:
 - A) If the branch depth attribute is defined (see 3.1.14) in the Snapshots Information attributes page of the source partition, then the branch depth attribute value for the destination partition shall be set to the same value as the branch depth attribute for the source partition; or
 - B) If the branch depth attribute is undefined (see 3.1.51) in the Snapshots Information attributes page of the source partition, then the branch depth attribute value for the destination partition shall be set to zero.

The destination snapshot partition shall be added as the newest entry in the history change as described in 6.30.5.

If the FREEZE bit is set to one, the device server shall:

- a) Note the value of the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the source partition for use in 6.10.4; and
- b) Set the object accessibility attribute in the Partition Information attributes page in the source partition to 0000 0001h.

The snapshot/clone tracking well known collection (see 4.6.6.5.3) shall be created in the destination partition, and initialized, including at least the following:

- a) Every user object and collection in the source partition shall have their User_Object_ID (see 4.6.5) or Collection_Object_ID (see 4.6.6) inserted as a member of the TRACKING collection (see 4.6.6.3); and
- b) The Command Tracking attributes page (see 7.1.3.20) shall be initialized to include at least the following:

- A) The percent complete attribute shall be set to zero;
- B) The active command status attribute shall be set to 88A9h (i.e., CREATE SNAPSHOT command in progress); and
- C) The ended command status attribute shall be set to FFFFh.

6.10.3 Processing after the IMMED_TR bit takes effect, if any

Every user object and collection in the source partition shall be duplicated in the destination snapshot partition using the:

- a) Duplication method (see 4.13.3) specified by the CDB; and
- b) Time of duplication method (see 4.13.4.2) specified by the CDB.

The membership and attributes of the snapshot/clone tracking well known collection for the destination partition should be maintained to restarting of an interrupted CREATE SNAPSHOT command with the minimum of repeated work (e.g., user objects or collections that have been fully duplicated should be removed from the snapshot/clone tracking well known collection). Other factors (e.g., meeting the requirements of the END time of duplication method (see 4.13.4.2)) may cause user objects and collections to be added to the snapshot/clone tracking well known collection.

6.10.4 Command completion

When an error is encountered or when all user objects and collections in the source partition have been duplicated in the destination snapshot partition as described in 6.10.3, the CREATE SNAPSHOT command processing shall be completed as described in this subclause.

If the FREEZE bit is set to one, the device server shall restore the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the source partition to the value noted in 6.10.2.

At least the following changes shall be made in the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the destination partition:

- a) The active command status attribute shall be set to zero;
- b) The ended command status attribute shall be set to indicate the condition (e.g., success or error) of the CREATE SNAPSHOT command processing; and
- c) If sense data is available, it shall be placed in the sense data attribute.

If the CREATE SNAPSHOT command processing complete (i.e., if the percent complete attribute in the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the destination partition is set to 100) and the ended command status attribute in the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the destination partition has been set to 0000h (i.e., GOOD status command completion), then the create completion time attribute in the Snapshots Information attributes page (see 7.1.3.30) in the destination snapshot partition shall be set to the value of the clock attribute in the Root Information attributes page (see 7.1.3.8); and

If the IMMED_TR bit is set to zero, status shall be returned for the CREATE SNAPSHOT command.

6.11 CREATE USER TRACKING COLLECTION

The CREATE USER TRACKING COLLECTION command (see table 87) creates a user tracking collection (see 3.1.55) and copies the membership of another collection of any type (see 4.6.6.1) to the newly created collection. If the specified output user tracking collection already exists and the contents of its Command Tracking attributes page (see 7.1.3.20) indicate that its use by another command has been finished, the membership of the output user tracking collection is replaced.

Table 90 — CREATE USER TRACKING COLLECTION command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8894h) _____ (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	PARTITION_ID (see 5.2.10) _____ (LSB)							
24	(MSB) _____							
31	REQUESTED COLLECTION_OBJECT_ID _____ (LSB)							
32	_____							
39	Reserved _____							
40	(MSB) _____							
47	SOURCE COLLECTION_OBJECT_ID _____ (LSB)							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field (see 5.2.10) specify the Partition_ID of the partition in which the user tracking collection is to be created.

The contents of the REQUESTED COLLECTION_OBJECT_ID field specify the Collection_Object_ID (see 4.6.6) to be assigned to the created user tracking collection. If the REQUESTED COLLECTION_OBJECT_ID field contains zero any Collection_Object_ID may be assigned. If the REQUESTED COLLECTION_OBJECT_ID field contains any value other than zero and the device server is unable to assign the requested Collection_Object_ID to the created user tracking collection, the user tracking collection shall not be created and the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

Within a partition, the device server shall not allow:

- a) The same Collection_Object_ID to be associated with more than one collection at any point in time; or
- b) A Collection_Object_ID to have the same value as any assigned User_Object_ID.

The contents of the SOURCE COLLECTION_OBJECT_ID field specify the Collection_Object_ID (see 4.6.6) of a collection that provides the initial membership of the created user tracking collection.

If the SOURCE COLLECTION_OBJECT_ID field contains zero, there shall be no members in the collection specified by the REQUESTED COLLECTION_OBJECT_ID field upon completion of the CREATE USER TRACKING COLLECTION command. This shall not be considered an error.

If the SOURCE COLLECTION_OBJECT_ID field does not contains zero, then:

- a) The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB if any of the following are true:
 - A) If the object specified by the SOURCE COLLECTION_OBJECT_ID field is not one of the following:
 - a) A LINKED collection (see 4.6.6.2);
 - b) A TRACKING collection (see 4.6.6.3); or
 - c) A SPONTANEOUS collection (see 4.6.6.2);or
 - B) A TRACKING collection is specified by the SOURCE COLLECTION_OBJECT_ID field and the active command status attribute in the Command Tracking attributes page (see 7.1.3.20) is not set to zero; and
- b) The membership of the collection specified by the SOURCE COLLECTION_OBJECT_ID field shall be copied to the collection specified by the REQUESTED COLLECTION_OBJECT_ID field.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB if any of the following are true:

- a) The SOURCE COLLECTION_OBJECT_ID field does not contain zero and the CDB CONTINUATION LENGTH field contains zero; or
- b) The SOURCE COLLECTION_OBJECT_ID field contains zero and the CDB CONTINUATION LENGTH field does not contain zero.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if:

- a) The CDB CONTINUATION LENGTH field does not contain zero; and
- b) The CDB continuation segment (see 5.3):
 - A) Does not contain one extension capabilities CDB continuation descriptor (see 5.4.6); or

- B) Contains any CDB continuation descriptors other than the extension capabilities CDB continuation descriptor.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15. The Collection_Object_ID assigned by the CREATE USER TRACKING COLLECTION command may be obtained from the Current Command attributes page (see 7.1.3.31).

The get and set attributes parameters shall affect only the created user tracking collection. The get and set attributes parameters shall not affect the attributes of the collection, if any, specified by the SOURCE COLLECTION_OBJECT_ID field.

The capability is described in 5.2.4. If the SOURCE COLLECTION_OBJECT_ID field does not contain zero, then the CREATE USER TRACKING COLLECTION command accesses two collections. One capability is necessary for each collection accessed. One capability appears in the CDB. The other capability appears in the CDB continuation segment (see 5.3). If the SOURCE COLLECTION_OBJECT_ID field contains zero, then the capability for the one collection accessed appears in the CDB.

The security parameters are described in 5.2.11.

The collection type attribute in the Collection Information attributes page (see 7.1.3.10) shall be set to 01h (i.e., TRACKING).

The Command Tracking attributes page (see 7.1.3.20) shall be initialized to include at least the following:

- a) The percent complete attribute shall be set to zero;
- b) The active command status attribute shall be set to zero (i.e., no command in progress); and
- c) The ended command status attribute shall be set to FFFFh (i.e., no ending status available).

The assigned Collection_Object_ID shall be placed in the Collection_Object_ID or User_Object_ID attribute in the Current Command attributes page (see 7.1.3.31).

If a CREATE USER TRACKING COLLECTION command causes the value in the number of collections, user tracking collections, and user objects attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the object count attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the object count quota.

If a CREATE USER TRACKING COLLECTION command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the object count quota.

6.12 DETACH CLONE

The DETACH CLONE command (see table 91) causes the OSD device server to change a clone partition into a primary partition (see 4.13.2).

Table 91 — DETACH CLONE command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (88AAh) _____ (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	CLONE PARTITION_ID _____ (LSB)							
24	_____							
47	Reserved _____							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The DETACH CLONE command accesses the following partitions:

- The clone partition that is specified by the CLONE PARTITION_ID field; and
- The source partition whose Partition_ID (see 4.6.4) is the value in the source partition attribute in the Snapshots Information attributes page (see 7.1.3.30) of the clone partition.

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The CLONE PARTITION_ID field contains the Partition_ID (see 4.6.4) of the clone partition that the DETACH CLONE command is being requested to detach.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB, if attributes in the Snapshots Information attributes page (see 7.1.3.30) of the clone partition have any of the following properties:

- a) The partition type attribute contains a value other than 02h (i.e., clone partition);
- b) The source partition attribute is undefined (see 3.1.51); or
- c) The create completion time attribute is undefined (see 3.1.51) and the refresh completion time attribute is undefined.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if the CDB continuation segment (see 5.3):

- a) Does not contain one extension capabilities CDB continuation descriptor (see 5.4.6); or
- b) Contains any CDB continuation descriptors other than the extension capabilities CDB continuation descriptor.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4. The DETACH CLONE command accesses two partitions. One capability is necessary for each partition accessed. One capability appears in the CDB. The other capability appears in the CDB continuation segment (see 5.3).

The security parameters are described in 5.2.11.

In the Snapshots Information attributes page (see 7.1.3.30) of the source partition, the following changes shall be made in attribute values:

- a) One shall be subtracted from the clones count attribute value; and
- b) The clone destination attribute whose value matches the Partition_ID (see 4.6.4) of the clone partition shall be made undefined (see 3.1.51).

In the Snapshots Information attributes page (see 7.1.3.30) of the clone partition, the following changes shall be made in attribute values:

- a) The partition type attribute shall be set to 00h (i.e., primary partition);
- b) The source partition attribute shall be made undefined (see 3.1.51);
- c) The branch depth attribute shall be set to zero;
- d) The create completion time attribute shall be made undefined;
- e) The refresh completion time attribute shall be made undefined.

The branch depth attributes in the Snapshots Information attributes page (see 7.1.3.30) of all partitions chained to the former clone partition shall be updated as follows:

- a) The branch depth attribute in all snapshot partitions that have the former clone partition as a source shall be set to zero;
- b) The branch depth attribute in each clone partition that has a source snapshot partition whose branch depth attribute is zero shall be set to one;

- c) The branch depth attribute in each snapshot partition that has a source clone partition whose branch depth is one shall be set to one;
- d) The branch depth attribute in each clone partition that has a source snapshot partition whose branch depth attribute is one shall be set to two;
- e) The branch depth attribute in each snapshot partition that has a source clone partition whose branch depth is two shall be set to two;
- f) The branch depth attribute in each clone partition that has a source snapshot partition whose branch depth attribute is n shall be set to n plus one; and
- g) The branch depth attribute in each snapshot partition that has a source clone partition whose branch depth is n plus one shall be set to n plus one.

6.13 FLUSH

The FLUSH command (see table 92) ensures that the specified data and attribute bytes for the specified user object are written to stable storage (see 4.14).

Table 92 — FLUSH command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8888h) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved		FLUSH SCOPE	
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	(MSB) _____							
31	USER_OBJECT_ID _____ (LSB)							
32	(MSB) _____							
39	FLUSH LENGTH _____ (LSB)							
40	(MSB) _____							
47	FLUSH STARTING BYTE ADDRESS _____ (LSB)							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The FLUSH SCOPE field (see table 93) specifies the scope of the data and attribute bytes that the device server shall ensure are written to stable storage.

Table 93 — User object flush scope values

Value	Scope of data and attributes written to stable storage	Range fields reserved ^a
00b	User object data and attributes	Yes
01b	User object attributes only	Yes
10b	User object data range and attributes	No
11b	Reserved	Yes
^a The range fields are the FLUSH LENGTH field and the FLUSH STARTING BYTE ADDRESS field.		

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

If the FLUSH SCOPE field contains 10b, the FLUSH LENGTH field specifies number of bytes that the device server shall ensure are written to stable storage.

If the FLUSH SCOPE field contains 10b, the FLUSH STARTING BYTE ADDRESS field specifies the location of the first byte of the flush length bytes that the device server shall ensure are written to stable storage relative to the first byte (i.e., byte zero) of the specified user object.

If the FLUSH SCOPE field contains 10b and the FLUSH STARTING BYTE ADDRESS field specifies a byte that is beyond the user object logical length attribute value in the User Object Information attributes page (see 7.1.3.11), then:

- a) No bytes shall be written to stable storage; and
- b) The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the FLUSH SCOPE field contains 10b, and the values in the FLUSH LENGTH field and FLUSH STARTING BYTE ADDRESS field result an attempt to write a byte that is beyond the user object logical length attribute value in the User Object Information attributes page to stable storage, then the device server shall ensure that the bytes between the flush starting byte address and the user object logical length are written to stable storage. This shall not be considered an error.

If the FLUSH SCOPE field contains 10b, an attempt to flush bytes that have never been written shall result in zeros being written to stable storage for those bytes. This shall not be considered an error.

The command data segment of the Data-Out Buffer is not used by the FLUSH command.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.14 FLUSH COLLECTION

The FLUSH COLLECTION command (see table 94) ensures that the specified collection information and attribute bytes for the specified collection are written to stable storage (see 4.14).

Table 94 — FLUSH COLLECTION command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (889Ah) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved		FLUSH SCOPE	
12	TIMESTAMPS CONTROL							
13	Reserved _____							
15								
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	(MSB) _____							
31	COLLECTION_OBJECT_ID _____ (LSB)							
32	Reserved _____							
47								
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	Get and set attributes parameters (see 5.2.6) _____							
79								
80	Capability (see 5.2.4) _____							
183								
184	Security parameters (see 5.2.11) _____							
235								

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The FLUSH SCOPE field (see table 95) specifies the scope of the collection information and attribute bytes that the device server shall ensure are written to stable storage.

Table 95 — Collection flush scope values

Value	Scope of data and attributes written to stable storage
00b	List of user objects contained in the collection
01b	Collection attributes only
10b	a) List of user objects contained in the collection; and b) Collection attributes
11b	Reserved

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The COLLECTION_OBJECT_ID field specifies Collection_Object_ID (see 4.6.6). If the collection identified by the COLLECTION_OBJECT_ID field does not exist, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The type of collection (see 4.6.6.1) being flushed shall not affect how the FLUSH COLLECTION command is processed (e.g., the attributes of the Command Tracking attributes page (see 7.1.3.20), if any, shall not be modified).

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.15 FLUSH OSD

The FLUSH OSD command (see table 96) ensures that the specified data and attribute bytes for the OSD logical unit are written to stable storage (see 4.14).

Table 96 — FLUSH OSD command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (889Ch) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved		FLUSH SCOPE	
12	TIMESTAMPS CONTROL							
13	Reserved _____							
47								
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	Get and set attributes parameters (see 5.2.6) _____							
79								
80	Capability (see 5.2.4) _____							
183								
184	Security parameters (see 5.2.11) _____							
235								

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The FLUSH SCOPE field (see table 97) specifies the scope of the data and attribute bytes that the device server shall ensure are written to stable storage.

Table 97 — Root object flush scope values

Value	Scope of data and attributes written to stable storage
00b	List of partitions contained in the OSD logical unit
01b	Root object attributes only
10b	a) List of partitions contained in the OSD logical unit; b) Root object attributes; c) Lists of user objects and collections contained in the every partition in the OSD logical unit; d) Partition attributes for every partition in the OSD logical unit; e) User object data for every user object in the OSD logical unit; f) User object attributes for every user object in the OSD logical unit; g) List of user objects contained in every the collection in the OSD logical unit; and h) Collection attributes for every collection in the OSD logical unit
11b	Reserved

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.16 FLUSH PARTITION

The FLUSH PARTITION command (see table 98) ensures that the specified data and attribute bytes for the specified user object are written to stable storage (see 4.14).

Table 98 — FLUSH PARTITION command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (889Bh) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved		FLUSH SCOPE	
12	TIMESTAMPS CONTROL							
13	Reserved _____							
15								
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	Reserved _____							
47								
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	Get and set attributes parameters (see 5.2.6) _____							
79								
80	Capability (see 5.2.4) _____							
183								
184	Security parameters (see 5.2.11) _____							
235								

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The FLUSH SCOPE field (see table 99) specifies the scope of the data and attribute bytes that the device server shall ensure are written to stable storage.

Table 99 — Partition flush scope values

Value	Scope of data and attributes written to stable storage
00b	List of user objects and collections contained in the partition
01b	Partition attributes only
10b	a) List of user objects and collections contained in the partition; b) Partition attributes; c) User object data for every user object in the partition; d) User object attributes for every user object in the partition; e) List of user objects contained in every the collection in the partition; and f) Collection attributes for every collection in the partition
11b	Reserved

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10. If the PARTITION_ID field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.17 FORMAT OSD

The FORMAT OSD command (see table 100) causes the OSD device server to delete all user objects, delete all partitions except partition zero, and set the attributes for the root object and partition zero as defined by this standard.

Table 100 — FORMAT OSD command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8881h) _____ (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved _____							
31								
32	(MSB) _____							
39	FORMATTED CAPACITY _____ (LSB)							
40	Reserved _____							
47								
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	Get and set attributes parameters (see 5.2.6) _____							
79								
80	Capability (see 5.2.4) _____							
183								
184	Security parameters (see 5.2.11) _____							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The FORMATTED CAPACITY field specifies the total capacity of the formatted OSD logical unit in bytes. If the FORMATTED CAPACITY field is set to zero, the entire logical unit is formatted as one OSD logical unit and the logical unit capacity established accordingly. If value in the FORMATTED CAPACITY field is greater than the maximum OSD logical unit capacity, the formatting command function shall process as if the FORMATTED CAPACITY field contained zero. This shall not be considered an error.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

During the processing of a FORMAT OSD command, the device server shall respond to commands as follows:

- a) In response to all commands except REQUEST SENSE and INQUIRY, the device server shall return CHECK CONDITION status unless a reservation conflict exists, in which case RESERVATION CONFLICT status shall be returned;
- b) In response to the INQUIRY command, the device server shall respond as specified in SPC-3; and
- c) In response to the REQUEST SENSE command, unless an error has occurred, the device server shall return GOOD status with parameter data containing the sense key set NOT READY and the additional sense code set to LOGICAL UNIT NOT READY FORMAT IN PROGRESS.

Upon successful completion of a FORMAT OSD command, the OSD logical unit shall contain:

- a) A root object;
- b) One partition OSD object for partition zero (see 3.1.33);
- c) Zero collections;
- d) Zero user objects;
- e) Root object attributes and partition zero attributes as defined by this standard;
- f) Vendor specific additional root object attributes and partition zero attributes;
- g) Root object attributes and partition zero attributes updated as specified by the CDB parameters;
- h) Zero collection attributes, if supported;
- i) Zero user object attributes;
- j) Zero attributes pages with page numbers between P + 1 0000h and P + 1FFF FFFFh; and
- k) Zero attributes pages with page numbers between R + 1 0000h and R + 1FFF FFFFh.

Processing of the FORMAT OSD command shall not alter the master key, root key, or any keys associated partition zero (see 4.12.9).

6.18 GET ATTRIBUTES

The GET ATTRIBUTES command (see table 101) instructs the device server to return the specified attributes for the specified root object, partition, collection, or user object before setting the attributes, if any, specified by the command (see 4.8.4).

Table 101 — GET ATTRIBUTES command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (888Eh) _____ (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved _____							
15								
16								
23	PARTITION_ID _____ (LSB)							
24	(MSB) _____							
31	USER_OBJECT_ID _____ (LSB)							
32	Reserved _____							
47								
48								
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	Get and set attributes parameters (see 5.2.6) _____							
79								
80								
183								
184	Security parameters (see 5.2.11) _____							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.19 GET MEMBER ATTRIBUTES

The GET MEMBER ATTRIBUTES command (see table 102) instructs the device server to return the specified attributes for the specified user tracking collection (see 3.1.55) and the user object members of the user tracking collection before setting the attributes, if any, specified by the command (see 4.8.4). The GET MEMBER ATTRIBUTES command is a multi-object command (see 4.6.6.6).

Table 102 — GET MEMBER ATTRIBUTES command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (88A2h) (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID (LSB)							
24	(MSB)							
31	COLLECTION_OBJECT_ID (LSB)							
32	Reserved							
47								
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5) (LSB)							
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6. Page format attribute processing is illegal for the GET MEMBER ATTRIBUTES command. If the GET/SET CDBFMT field contains a value other than 11b, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The COLLECTION_OBJECT_ID field specifies Collection_Object_ID (see 4.6.6) to be processed. The device server shall constrain the Collection_Object_ID values as described in 4.6.6.6.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15. Get and set attributes processing requirements specific to multi-object commands are described in 4.6.6.6.

The same attributes (i.e., the same combinations of attribute page and attribute number) are retrieved and/or set for all user objects in the specified collection. If user object attributes are set, all such attributes are set to the same values in all user objects in the specified collection.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.20 LIST

6.20.1 Introduction

The LIST command is used to obtain information from the root object or a partition.

Table 103 — LIST command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8883h) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved	LIST_ATTR	GET/SET CDBFMT		SORT ORDER			
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	_____							
31	Reserved _____							
32	(MSB) _____							
39	ALLOCATION LENGTH _____ (LSB)							
40	(MSB) _____							
47	INITIAL OBJECT_ID _____ (LSB)							
48	(MSB) _____							
51	LIST IDENTIFIER _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The contents of the ISOLATION field are described in 5.2.8.

The LIST_ATTR bit value combined with the value in the PARTITION_ID field (see table 104) specify the information that shall be returned.

Table 104 — Specifying objects and attributes to be listed

PARTITION_ID field	LIST_ATTR field	Description
zero	zero	The Partition_IDs (see 4.6.4) in the root object shall be returned in the parameter data. Partition_ID zero shall not be returned in the parameter data, but any requested attributes shall be returned in the retrieved attributes segment of the Data-In Buffer (see 4.15.3).
	one	The Partition_IDs in the root object and attributes specified by the get and set attributes parameters (see 5.2.6) for each partition shall be returned in the parameter data. Partition_ID zero shall not be returned in the parameter data, but any requested attributes shall be returned in the retrieved attributes segment of the Data-In Buffer.
non-zero	zero	The User_Object_IDs in the specified partition shall be returned in the parameter data.
	one	The User_Object_IDs in the specified partition and attributes specified by the get and set attributes parameters for each user object shall be returned in the parameter data.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6. If the LIST_ATTR bit is set to one, page format attribute processing is illegal. If the LIST_ATTR bit is set to one and the GET/SET CDBFMT field contains a value other than 11b, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The SORT ORDER field (see table 105) specifies the order in which the returned Partition_IDs or User_Object_IDs shall be sorted.

Table 105 — LIST command sort order values

Sort Order	Description
0h	Ascending numeric value
1h to Fh	Reserved

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10. The contents of the PARTITION_ID field combined with the LIST_ATTR bit value specify the information that shall be returned (see table 104).

The LIST IDENTIFIER field contains zero if the INITIAL OBJECT_ID field contains Partition_ID or User_Object_ID (see 4.6.2). Otherwise, the LIST IDENTIFIER field contains the list identifier returned by a previous LIST command.

The ALLOCATION LENGTH field is described in 5.2.2.

The contents of the INITIAL OBJECT_ID field depend on the contents of the LIST IDENTIFIER field. If the LIST IDENTIFIER field contains zero, the INITIAL OBJECT_ID field specifies the lowest valued Partition_ID or User_Object_ID to be returned. If the LIST IDENTIFIER field contains any value other than zero, the INITIAL OBJECT_ID field contains the value in the CONTINUATION OBJECT_ID field from the same returned parameter data that contained the value in the LIST IDENTIFIER field.

If the LIST_ATTR bit is set to zero, the get attributes parameters are described in 5.2.6. If the LIST_ATTR bit is set to one, the get attributes parameters are described in 5.2.6.4 with the additional processing requirements that depend on PARTITION_ID field contents and attribute page numbers as shown in table 106.

Table 106 — Attributes processing requirements for LIST commands with the LIST_ATTR bit set to one

PARTITION_ID field	Attribute page number values	Description
zero	R+0h to R+2FFF FFFFh	The retrieved attribute values shall be returned in the retrieved attributes segment of the Data-In Buffer (see 4.15.3) in the format described in 7.1.4.
	P+0h to P+2FFF FFFFh	a) For Partition_ID zero, the retrieved attribute values shall be returned in the retrieved attributes segment of the Data-In Buffer in the format described in 7.1.4; and b) For any Partition_ID other than zero, the retrieved attribute values shall be returned in the LIST command parameter data as described in 6.20.2.
	FFFF FFFEh	The retrieved Current Command attributes page (see 7.1.3.31) attribute values shall be returned in the retrieved attributes segment of the Data-In Buffer in the format described in 7.1.4.
	All other values	The command to be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.
non-zero	P+0h to P+2FFF FFFFh	The retrieved attribute values shall be returned in the retrieved attributes segment of the Data-In Buffer in the format described in 7.1.4.
	0h to 2FFF FFFFh	The retrieved attribute values shall be returned in the LIST command parameter data as described in 6.20.2.
	FFFF FFFEh	The retrieved Current Command attributes page (see 7.1.3.31) attribute values shall be returned in the retrieved attributes segment of the Data-In Buffer in the format described in 7.1.4.
	All other values	The command to be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

Regardless of LIST_ATTR bit value, the set attributes parameters are described in 5.2.6.

The format of the Data-Out Buffer when attributes are being retrieved or set is described in 4.15.4.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.20.2 LIST command parameter data

The parameter data returned by the LIST command (see table 107) contains the requested list of partitions or user objects.

Table 107 — LIST command parameter data

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
7	ADDITIONAL LENGTH (n-7) _____ (LSB)							
8	(MSB) _____							
15	CONTINUATION OBJECT_ID _____ (LSB)							
16	(MSB) _____							
19	LIST IDENTIFIER _____ (LSB)							
20	Reserved _____							
22								
23	OBJECT DESCRIPTOR FORMAT						LSTCHG	Obsolete
	Object descriptor list							
24	Object descriptor [first] _____							
	⋮							
n	Object descriptor [last] _____							

The ADDITIONAL LENGTH field indicates the number of bytes of LIST command parameter data that follow. If the parameter data is truncated due to insufficient allocation length, the ADDITIONAL LENGTH field shall not be altered to reflect the truncation (i.e., the additional length indicates the number of bytes that would follow if the allocation length had been infinite). If the untruncated number of bytes that follow is greater than FFFF FFFF FFFF FFFFh the additional length shall be set to FFFF FFFF FFFF FFFFh.

The CONTINUATION OBJECT_ID field provides information that may be used to continue a truncated list with a new LIST command. If the CONTINUATION OBJECT_ID field contains zero, the parameter data contains all of the list results and no further LIST commands are needed. If a new LIST command is sent to continue a truncated list, the contents of the CONTINUATION OBJECT_ID field are copied to the INITIAL OBJECT_ID field of that new command.

The LIST IDENTIFIER field contains an identifier required for continuing a truncated list in a new LIST command. If a new LIST command is sent to continue a truncated list, the contents of the LIST IDENTIFIER field are copied to the LIST IDENTIFIER field of that new command.

The OBJECT DESCRIPTOR FORMAT field (see table 108) indicates the format of the object descriptors.

Table 108 — LIST command OBJECT DESCRIPTOR FORMAT field values

Code	Description	Reference
00h	The LIST command parameter data format shall be as specified in OSD (see 2.2).	
01h	The object descriptors are a list of Partition_IDs each of which is eight bytes and has the format shown in table 109.	6.20.3.1
02h	Each object descriptor has a multi-object retrieved attributes format (see table 230), and contains a Partition_ID followed by attribute parameters associated with the indicated partition.	7.1.4.4
03h to 20h	Reserved	
21h	The object descriptors are a list of User_Object_IDs each of which is eight bytes and has the format shown in table 111.	6.20.3.3
22h	Each object descriptor has a multi-object retrieved attributes format (see table 230), and contains a User_Object_ID followed by attribute parameters associated with the indicated user object.	7.1.4.4
23h to 3Fh	Reserved	

A LSTCHG (list has changed) bit set to zero indicates that the entries in the list of OSD objects in the parameter data has not changed since the first LIST command identified by the list identifier. A LSTCHG bit set to one indicates that the entries in the list of OSD objects in the parameter data has changed since the first LIST command identified by the list identifier and that starting the list over at the original initial object_id may be necessary in order to obtain a complete list.

The parameter data shall contain one object descriptor for each user object or partition identified by the LIST command. If the list is truncated based on allocation length, the truncation shall not occur in the middle of an object descriptor.

The LIST command parameter data shall not contain Collection_Object_IDs. Lists of Collection_Object_IDs may be obtained using the LIST COLLECTION command (see 6.21).

6.20.3 LIST command and LIST COLLECTION command object descriptor formats

6.20.3.1 Partition_ID only object descriptor format

For a LIST command with the PARTITION_ID field set to zero and the LIST_ATTR bit set to zero, each parameter data object descriptor shall be eight bytes in length and shall have the format shown in table 109.

Table 109 — Partition_ID only object descriptor format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
7	PARTITION_ID							(LSB)

The PARTITION_ID field indicates the partition (see 4.6.4) to which the object descriptor applies.

6.20.3.2 Collection_Object_ID only object descriptor format

For a LIST COLLECTION command with the COLLECTION_OBJECT_ID field set to zero and the LIST_ATTR bit set to zero, each parameter data object descriptor shall be eight bytes in length and shall have the format shown in table 110.

Table 110 — Collection_Object_ID only object descriptor format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
7	COLLECTION_OBJECT_ID _____ (LSB)							

The COLLECTION_OBJECT_ID field indicates the collection (see 4.6.6) to which the object descriptor applies.

6.20.3.3 User_Object_ID only object descriptor format

Each parameter data object descriptor shall be eight bytes in length and shall have the format shown in table 111 if the LIST_ATTR bit is set to zero for:

- a) A LIST command with the PARTITION_ID field set to a value other than zero; or
- b) A LIST COLLECTION command with the COLLECTION_OBJECT_ID field set to a value other than zero.

Table 111 — User_Object_ID only object descriptor format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
7	USER_OBJECT_ID (LSB)							

The USER_OBJECT_ID field indicates the user object (see 4.6.5) to which the object descriptor applies.

6.21 LIST COLLECTION

The LIST COLLECTION command (see table 112) is used to get information from a collection.

Table 112 — LIST COLLECTION command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8897h) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved	LIST_ATTR	GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	(MSB) _____							
31	COLLECTION_OBJECT_ID _____ (LSB)							
32	(MSB) _____							
39	ALLOCATION LENGTH _____ (LSB)							
40	(MSB) _____							
47	INITIAL_OBJECT_ID _____ (LSB)							
48	(MSB) _____							
51	LIST_IDENTIFIER _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The contents of the ISOLATION field are described in 5.2.8.

The LIST_ATTR bit value combined with the value in the COLLECTION_OBJECT_ID field (see table 113) specify the information that shall be returned.

Table 113 — Specifying collections and attributes to be listed

COLLECTION_OBJECT_ID field	LIST_ATTR field	Description
zero	zero	The Collection_Object_IDs (see 4.6.6) in the specified partition shall be returned in the parameter data.
	one	The Collection_Object_IDs in the specified partition and attributes specified by the get and set attributes parameters (see 5.2.6) for each collection shall be returned in the parameter data.
non-zero	zero	The User_Object_IDs in the specified collection shall be returned in the parameter data.
	one	The User_Object_IDs in the specified collection and attributes specified by the get and set attributes parameters for each user object shall be returned in the parameter data.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6. If the LIST_ATTR bit is set to one, page format attribute processing is illegal. If the LIST_ATTR bit is set to one and the GET/SET CDBFMT field contains a value other than 11b, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The COLLECTION_OBJECT_ID field specifies Collection_Object_ID (see 4.6.6) to be processed. The contents of the COLLECTION_OBJECT_ID field combined with the LIST_ATTR bit value specify the information that shall be returned (see table 113). If the collection identified by a non-zero COLLECTION_OBJECT_ID field does not exist, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The type of collection (see 4.6.6.1) being flushed shall not affect how the LIST COLLECTION command is processed (e.g., the attributes of the Command Tracking attributes page (see 7.1.3.20), if any, shall not be modified).

The LIST IDENTIFIER field contains zero if the INITIAL_OBJECT_ID field contains Collection_Object_ID or User_Object_ID (see 4.6.5). Otherwise, the LIST IDENTIFIER field contains the list identifier returned by a previous LIST COLLECTION command.

The ALLOCATION LENGTH field is described in 5.2.2.

The contents of the INITIAL_OBJECT_ID field depend on the contents of the LIST IDENTIFIER field. If the LIST IDENTIFIER field contains zero, the INITIAL_OBJECT_ID field specifies the lowest valued Collection_Object_ID or User_Object_ID to be returned. If the LIST IDENTIFIER field contains any value other than zero, the INITIAL_OBJECT_ID field contains the value in the CONTINUATION_OBJECT_ID field from the same returned parameter data that contained the value in the LIST IDENTIFIER field.

If the LIST_ATTR bit is set to zero, the get attributes parameters are described in 5.2.6. If the LIST_ATTR bit is set to one, the get attributes parameters are described in 5.2.6.4 with the additional processing requirements that depend on COLLECTION_OBJECT_ID field contents and attribute page numbers as shown in table 114.

Table 114 — Attributes processing requirements for LIST COLLECTION commands with the LIST_ATTR bit set to one

COLLECTION_OBJECT_ID field	Attribute page number values	Description
zero	P+0h to P+2FFF FFFFh	The retrieved attribute values shall be returned in the retrieved attributes segment of the Data-In Buffer (see 4.15.3) in the format described in 7.1.4.
	C+0h to C+2FFF FFFFh	The retrieved attribute values shall be returned in the LIST COLLECTION command parameter data as described in 6.20.2.
	FFFF FFFEh	The retrieved Current Command attributes page (see 7.1.3.31) attribute values shall be returned in the retrieved attributes segment of the Data-In Buffer in the format described in 7.1.4.
	All other values	The command to be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.
non-zero	C+0h to C+2FFF FFFFh	The retrieved attribute values shall be returned in the retrieved attributes segment of the Data-In Buffer in the format described in 7.1.4.
	0h to 2FFF FFFFh	The retrieved attribute values shall be returned in the LIST COLLECTION command parameter data as described in 6.20.2.
	FFFF FFFEh	The retrieved Current Command attributes page (see 7.1.3.31) attribute values shall be returned in the retrieved attributes segment of the Data-In Buffer in the format described in 7.1.4.
	All other values	The command to be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

Regardless of LIST_ATTR bit value, the set attributes parameters are described in 5.2.6.

The format of the Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

The parameter data returned by the LIST COLLECTION command (see table 115) contains the requested information about the collections in the specified partition or user objects in the specified collection.

Table 115 — LIST COLLECTION command parameter data

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	ADDITIONAL LENGTH (n-7)						(LSB)
7								
8	(MSB)	CONTINUATION OBJECT_ID						(LSB)
15								
16	(MSB)	LIST IDENTIFIER						(LSB)
19								
20		Reserved						
22								
23	OBJECT DESCRIPTOR FORMAT						LSTCHG	Obsolete
	Object descriptor list							
24		Object descriptor [first]						
	⋮							
n		Object descriptor [last]						

The ADDITIONAL LENGTH field indicates the number of bytes of LIST COLLECTION command parameter data that follow. If the parameter data is truncated due to insufficient allocation length, the ADDITIONAL LENGTH field shall not be altered to reflect the truncation (i.e., the additional length indicates the number of bytes that would follow if the allocation length had been infinite). If the untruncated number of bytes that follow is greater than FFFF FFFF FFFF FFFFh the additional length shall be set to FFFF FFFF FFFF FFFFh.

The CONTINUATION OBJECT_ID field provides information that may be used to continue a truncated list with a new LIST COLLECTION command. If the CONTINUATION OBJECT_ID field contains zero, the parameter data contains all of the list results and no further LIST COLLECTION commands are needed. If a new LIST COLLECTION command is sent to continue a truncated list, the contents of the CONTINUATION OBJECT_ID field are copied to the INITIAL OBJECT_ID field of that new command.

The LIST IDENTIFIER field contains an identifier required for continuing a truncated list in a new LIST COLLECTION command. If a new LIST COLLECTION command is sent to continue a truncated list, the contents of the LIST IDENTIFIER field are copied to the LIST IDENTIFIER field of that new command.

The OBJECT DESCRIPTOR FORMAT field (see table 116) indicates the format of the object descriptors.

Table 116 — LIST COLLECTION command OBJECT DESCRIPTOR FORMAT field values

Code	Description	Reference
00h	The LIST COLLECTION command parameter data format shall be as specified in OSD (see 2.2).	
01h to 10h	Reserved	
11h	The object descriptors are a list of Collection_Object_IDs each of which is eight bytes and has the format shown in table 110.	6.20.3.2
12h	Each object descriptor has a multi-object retrieved attributes format (see table 230), and contains a Collection_Object_ID followed by attribute parameters associated with the indicated collection.	7.1.4.4
13h to 20h	Reserved	
21h	The object descriptors are a list of User_Object_IDs each of which is eight bytes and has the format shown in table 111.	6.20.3.3
22h	Each object descriptor has a multi-object retrieved attributes format (see table 230), and contains a User_Object_ID followed by attribute parameters associated with the indicated user object.	7.1.4.4
23h to 3Fh	Reserved	

A LSTCHG (list has changed) bit set to zero indicates that the entries in the list of OSD objects in the parameter data has not changed since the first LIST COLLECTION command identified by the list identifier. A LSTCHG bit set to one indicates that the entries in the list of OSD objects in the parameter data has changed since the first LIST COLLECTION command identified by the list identifier and that starting the list over at the original initial object_id may be necessary in order to obtain a complete list.

The parameter data shall contain one object descriptor for each user object or collection identified by the LIST COLLECTION command. If the list is truncated based on allocation length, the truncation shall not occur in the middle of an object descriptor.

6.22 OBJECT STRUCTURE CHECK

6.22.1 Introduction

The OBJECT STRUCTURE CHECK command (see table 117) verifies the integrity of the OBSD storage for a partition or for the root object and all partitions.

Table 117 — OBJECT STRUCTURE CHECK command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (8880h)							
10	(LSB)							
10	Reserved							
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID TO CHECK							
23	(LSB)							
24	Reserved							
47								
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
51	(LSB)							
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID TO CHECK field specify the Partition_ID (see 4.6.4) for which the structure checking operation shall be performed. If the PARTITION_ID TO CHECK field contains zero, the structure checking shall be performed on the root object and all partitions. If the non-zero partition identified by the PARTITION_ID TO CHECK field does not exist, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

Upon completion of an OBJECT STRUCTURE COMMAND with GOOD status:

- a) The OBSD shall be ready to process application client commands to the addressed partition or to the root object without need for the processing of additional OBJECT STRUCTURE CHECK commands until a hard reset SCSI device condition is established in response to an event (see SAM-4); and
- b) Any uncorrectable storage damage detected by processing the OBJECT STRUCTURE CHECK command shall be reported as described in 4.11.3.1.

The detection of uncorrectable storage damage shall not cause an OBJECT STRUCTURE CHECK command to be terminated with CHECK CONDITION status.

Following completion of an OBJECT STRUCTURE CHECK command, the application client should send a REQUEST SENSE command to retrieve the unit attention condition, if any, established as a result of updating the damage storage reporting information as described in 4.11.3.1.

6.22.2 Structure checking for the root object and all partitions

Before modifying any OBSD storage in response to an OBJECT STRUCTURE CHECK command for the root object and all partitions, the device server shall ensure that all commands received prior to the OBJECT STRUCTURE CHECK command have completed processing.

While an OBJECT STRUCTURE CHECK command for the root object and all partitions is being processed, the device server shall terminate all commands except an INQUIRY command (see SPC-4), a REPORT LUNS command (see SPC-4), or a REQUEST SENSE command (see SPC-4) as follows:

- a) CHECK CONDITION status;
- b) The sense key set to NOT READY;
- c) The additional sense code set to LOGICAL UNIT NOT READY, REBUILD IN PROGRESS;
- d) The INFORMATION field set to zero;
- e) The SKSV bit set to one; and
- f) The sense key specific data containing a progress indication as described in SPC-4.

While an OBJECT STRUCTURE CHECK command for the root object and all partitions is being processed, the device server shall complete REQUEST SENSE commands with GOOD status and the following sense data:

- a) The sense key set to NOT READY;
- b) The additional sense code set to LOGICAL UNIT NOT READY, REBUILD IN PROGRESS;
- c) The INFORMATION field set to zero;
- d) The SKSV bit set to one; and
- e) The sense key specific data containing a progress indication as described in SPC-4.

The processing of an OBJECT STRUCTURE CHECK command shall not affect the processing of the INQUIRY command or a REPORT LUNS command.

6.22.3 Structure checking for a specific partition

Before modifying any OBSD storage associated with the specified partition in response to an OBJECT STRUCTURE CHECK command for a specific partition, the device server shall ensure that all commands addressed to the specified partition received prior to the OBJECT STRUCTURE CHECK command have completed processing.

While an OBJECT STRUCTURE CHECK command for a specified partition is being processed, the device server shall terminate all commands addressed to that partition with:

- a) CHECK CONDITION status;
- b) The sense key set to NOT READY;
- c) The additional sense code set to LOGICAL UNIT NOT READY, REBUILD IN PROGRESS;
- d) The INFORMATION field set to the Partition_ID of a partition being processed by an OBJECT STRUCTURE CHECK command;
- e) The SKSV bit set to one; and
- f) The sense key specific data containing a progress indication as described in SPC-4.

6.23 PERFORM SCSI COMMAND

The PERFORM SCSI COMMAND command (see table 118) allows an implemented SPC-3 command (e.g., LOG SENSE) to be processed when the security method is not NOSEC (see 4.12.4). The PERFORM SCSI COMMAND command also allows an implemented SPC-4 command to be processed concurrently with attributes retrieval and setting command functions.

Table 118 — PERFORM SCSI COMMAND command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8F7Ch) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	PARTITION_ID (0000 0000 0000 0000h) _____ (LSB)							
24	(MSB) _____							
31	USER_OBJECT_ID (0000 0000 0000 0000h) _____ (LSB)							
32	_____							
47	REQUEST CDB _____							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10. Because the PERFORM SCSI COMMAND affects the OSD logical unit, it is addressed to the root object (i.e., Partition_ID zero). If the PARTITION_ID field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the USER_OBJECT_ID field are described in 5.2.14. Because the PERFORM SCSI COMMAND affects the OSD logical unit, it is addressed to the root object (i.e., User_Object_ID zero). If the USER_OBJECT_ID field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The REQUEST CDB field contains the fixed-length CDB for the SPC-4 command to be processed. Any bytes between the end of the CDB for the SPC-4 command and the end of the REQUEST CDB field shall be ignored (e.g., a ten-byte CDB occupies the first ten bytes of the REQUEST CDB field and the remaining six bytes are ignored).

If SPC-4 command specified by the REQUEST CDB field is not one of the commands listed in table 119, the PERFORM SCSI COMMAND command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

Table 119 — Request CDBs allowed in the PERFORM SCSI COMMAND

Command name	Operation code	Service action (if any)
INQUIRY	12h	
LOG SELECT	4Ch	
LOG SENSE	4Dh	
MODE SELECT(10)	55h	
MODE SENSE(10)	5Ah	
READ BUFFER	3Ch	
RECEIVE DIAGNOSTIC RESULTS	1Ch	
REPORT LUNS	A0h	
REPORT SUPPORTED OPERATION CODES	A3h	0Ch
REPORT SUPPORTED TASK MANAGEMENT FUNCTIONS	A3h	0Dh
REPORT TARGET PORT GROUPS	A3h	0Ah
REQUEST SENSE	03h	
SEND DIAGNOSTIC	1Dh	
SET TARGET PORT GROUPS	A4h	0Ah
TEST UNIT READY	00h	
WRITE BUFFER	3Bh	

Only those commands specified as mandatory to implement in table 80 (see 6.1) are mandatory to implement in this command. If the SPC-4 command specified by the REQUEST CDB field is not implemented, the PERFORM SCSI COMMAND command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the SPC-4 command specified by the REQUEST CDB field transfers data to the Data-In Buffer, the data bytes shall be placed in the traditional command or parameter data segment of the Data-In Buffer (see 4.15). If the SPC-4 command specified by the REQUEST CDB field transfers data from the Data-Out Buffer, the data bytes shall be retrieved from the traditional command or parameter data segment of the Data-Out Buffer.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

| The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

If the PERFORM SCSI COMMAND is terminated with CHECK CONDITION status, the sense key is ILLEGAL REQUEST, the sense key specific sense data descriptor (see SPC-4) is included in the sense data, and the C/D bit is set to one, then values in the FIELD POINTER field shall be based on the PERFORM SCSI COMMAND CDB (i.e., not on the CDB in the REQUEST CDB field).

6.24 PERFORM TASK MANAGEMENT FUNCTION

The PERFORM TASK MANAGEMENT FUNCTION command (see table 120) allows a SAM-4 task management function (e.g., ABORT TASK) to be processed when the security method is not NOSEC (see 4.12.4). The PERFORM TASK MANAGEMENT FUNCTION command also allows a SAM-4 task management function to be processed concurrently with attributes retrieval and setting command functions.

Table 120 — PERFORM TASK MANAGEMENT FUNCTION command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (8F7Dh)							
10	(LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13								
15	Reserved							
16	(MSB)							
23	PARTITION_ID							
23	(LSB)							
24	(MSB)							
31	USER_OBJECT_ID							
31	(LSB)							
32	(MSB)							
33	ALLOCATION LENGTH							
33	(LSB)							
34								
38	Reserved							
39	TASK MANAGEMENT FUNCTION							
40								
47	TASK TAG							
47								
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
51	(LSB)							
52								
79	Get and set attributes parameters (see 5.2.6)							
80								
183	Capability (see 5.2.4)							
183								
184								
235	Security parameters (see 5.2.11)							
235								

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

The ALLOCATION LENGTH field is described in 5.2.2.

The TASK MANAGEMENT FUNCTION field (see table 121) specifies the SAM-4 task management function to be processed.

Table 121 — TASK MANAGEMENT FUNCTION field

Code	SAM-4 Task Management Function	Addressed OSD Object	Task Tag Specified	ADDITIONAL RESPONSE INFORMATION parameter data field reserved
01h	ABORT TASK	Any	Yes	Yes
02h	ABORT TASK SET	Root	No	Yes
04h	CLEAR TASK SET	Root	No	Yes
08h	LOGICAL UNIT RESET	Root	No	Yes
10h	I_T NEXUS RESET	Root	No	Yes
40h	CLEAR ACA	Any	No	Yes
80h	QUERY TASK	Any	Yes	Yes
81h	QUERY TASK SET	Any	Yes	Yes
82h	QUERY UNIT ATTENTION	Any	Yes	No
All codes not listed in this table are reserved.				

If the TASK MANAGEMENT FUNCTION field contains a value that table 121 lists as being addressed to the root object and either the PARTITION_ID field or the USER_OBJECT_ID field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The TASK TAG field contains the task tag that identifies the task to be managed if the TASK MANAGEMENT FUNCTION field contains a value listed as specifying a task tag in table 121. If table 121 lists a task management function as not specifying a task tag, then the contents of the task tag field shall be ignored.

The format of the task tag is specified in the applicable SCSI transport protocol standard and the length of the task tag may be less than eight bytes. Any bytes between the end of the task tag and the end of the TASK TAG field shall be ignored (e.g., a two-byte task tag occupies the first two bytes of the TASK TAG field and the remaining six bytes are ignored).

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

The PERFORM TASK MANAGEMENT FUNCTION parameter data (see table 122) indicates the results of the task management function specified by the TASK MANAGEMENT FUNCTION field. The parameter data shall be returned in the command data or parameter data segment of the Data-In Buffer (see 4.15.3).

Table 122 — PERFORM TASK MANAGEMENT FUNCTION parameter data format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
1	ADDITIONAL LENGTH (0006h)							(LSB)
2	Reserved							
3								
4	SERVICE RESPONSE							
5	(MSB)							
7	ADDITIONAL RESPONSE INFORMATION							(LSB)

The ADDITIONAL LENGTH field indicates the number of bytes that follow. The ADDITIONAL LENGTH field shall contain six.

The SERVICE RESPONSE field (see table 123) indicates the service response (see SAM-4) returned by the task management function specified by the TASK MANAGEMENT FUNCTION field.

Table 123 — SERVICE RESPONSE field

Code	Service response (see SAM-4)
00h	FUNCTION COMPLETE
05h	FUNCTION SUCCEEDED
08h	FUNCTION REJECTED
09h	INCORRECT LOGICAL UNIT NUMBER
All codes not listed in this table are reserved.	

If the TASK MANAGEMENT FUNCTION field in the CDB specifies a task management function for which the contents of the ADDITIONAL RESPONSE INFORMATION field are not reserved (see table 122), then the ADDITIONAL RESPONSE INFORMATION field shall contain the data defined by SAM-4 for the Additional Response Information argument for the specified task management function.

6.25 PUNCH

The PUNCH command (see table 124) removes bytes from a user object.

Table 124 — PUNCH command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8884h) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	(MSB) _____							
31	USER_OBJECT_ID _____ (LSB)							
32	(MSB) _____							
39	PUNCH LENGTH _____ (LSB)							
40	(MSB) _____							
47	PUNCH STARTING BYTE ADDRESS _____ (LSB)							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

The PUNCH LENGTH field specifies number of bytes to be removed from the user object. A punch length of zero shall cause no bytes to be removed from the user object. This shall not be considered an error.

The PUNCH STARTING BYTE ADDRESS field specifies the location where the removal of bytes from the specified user object is to commence relative to the first byte (i.e., byte zero) of the user object (e.g., if the punch starting byte address is five and the punch length is two, then byte seven in the user object becomes byte five, and so on for the remaining logical length of the user object).

If the PUNCH STARTING BYTE ADDRESS field specifies a byte that is beyond the user object logical length attribute value in the User Object Information attributes page (see 7.1.3.11), then:

- a) No bytes shall be removed from the user object; and
- b) The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the values in the PUNCH LENGTH field and PUNCH STARTING BYTE ADDRESS field result an attempt to remove bytes that are beyond the user object logical length attribute value in the User Object Information attributes page, then the user object shall be truncated by setting the user object logical length attribute value in the User Object Information attributes page to the value in the PUNCH STARTING BYTE ADDRESS field. This shall not be considered an error.

The command data segment of the Data-Out Buffer is not used by the PUNCH command.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

If a PUNCH command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated. The quota testing principles described in 4.10.3 apply to the testing of the maximum capacity quota.

6.26 QUERY

6.26.1 Introduction

The QUERY command (see table 125) instructs the device server to return a list of the user objects that are members of the specified user tracking collection (see 3.1.55) and have attributes matching the specified values. The QUERY command is a multi-object command (see 4.6.6.6).

Table 125 — QUERY command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (88A0h) _____ (LSB)							
10	Reserved					ISOLATION		
11	IMMED_TR	Reserved	GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	(MSB) _____							
31	COLLECTION_OBJECT_ID _____ (LSB)							
32	(MSB) _____							
39	ALLOCATION LENGTH _____ (LSB)							
40	(MSB) _____							
47	MATCHES COLLECTION_OBJECT_ID _____ (LSB)							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The contents of the ISOLATION field are described in 5.2.8.

The IMMED_TR bit is described in 5.2.5. If the IMMED_TR bit is set to one, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB if any of the following are true:

- a) The ALLOCATION LENGTH field is not set to zero; or
- b) The MATCHES COLLECTION_OBJECT_ID field is set to zero.

If the IMMED_TR bit is set to one, the user tracking collections (see 3.1.55) specified by the COLLECTION_OBJECT_ID field and the MATCHES COLLECTION_OBJECT_ID field shall be initialized before GOOD status is returned.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The COLLECTION_OBJECT_ID field specifies Collection_Object_ID (see 4.6.6) of the user tracking collection (see 3.1.55) to be processed. The device server shall constrain the Collection_Object_ID values as described in 4.6.6.6.

The ALLOCATION LENGTH field is described in 5.2.2 and specifies the allocation length for the matches list parameter data (see 6.26.2).

If the ALLOCATION LENGTH field contains a non-zero value and the MATCHES COLLECTION_OBJECT_ID field contains a non-zero value, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The MATCHES COLLECTION_OBJECT_ID field specifies Collection_Object_ID (see 4.6.6) of the user tracking collection (see 3.1.55) in which the QUERY command shall place the User_Object_IDs of the user objects that are members of the collection specified by the COLLECTION_OBJECT_ID field and have attributes matching the specified values. If the MATCHES COLLECTION_OBJECT_ID field is set to zero, the matching User_Object_IDs shall not be returned in a user tracking collection.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB if any of the following are true:

- a) The contents of the COLLECTION_OBJECT_ID field are identical to the contents of the MATCHES COLLECTION_OBJECT_ID field;
- b) The COLLECTION_OBJECT_ID field contains a value that is not the Collection_Object_ID of a user tracking collection (see 3.1.55);
- c) The active command status attribute in the Command Tracking attributes page (see 7.1.3.20) of the user tracking collection specified by the COLLECTION_OBJECT_ID field contains a value other than 0000h.
- d) The MATCHES COLLECTION_OBJECT_ID field contains a non-zero value that is not the Collection_Object_ID of a user tracking collection; or
- e) The active command status attribute in the Command Tracking attributes page (see 7.1.3.20) of the user tracking collection specified by the MATCHES COLLECTION_OBJECT_ID field contains a value other than 0000h.

If the matches user tracking collection contains any members, they shall be removed before the user tracking collection specified by the COLLECTION_OBJECT_ID field is processed.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if CDB continuation segment (see 5.3):

- a) Does not contain one query list CDB continuation descriptor (see 5.4.3);
- b) Contains more than one extension capabilities CDB continuation descriptor (see 5.4.6) if the MATCHES COLLECTION_OBJECT_ID field contains a non-zero value;

- c) Contains more than zero extension capabilities CDB continuation descriptor (see 5.4.6) if the MATCHES COLLECTION_OBJECT_ID field contains zero;
- d) Contains any CDB continuation descriptors other than the following:
 - A) Query list CDB continuation descriptor (see 5.4.3); and
 - B) Extension capabilities CDB continuation descriptor (see 5.4.6).

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15. The get and set attributes processing requirements specific to multi-object commands (see 4.6.6.6) shall affect only the collection specified by the COLLECTION_OBJECT_ID field. The get and set attributes parameters shall not affect the attributes of the collection specified by the MATCHES COLLECTION_OBJECT_ID field.

The capability is described in 5.2.4. The QUERY command may access two user tracking collections. One capability is necessary for each user tracking collection accessed. If two user tracking collections are accessed, one capability appears in the CDB, and the other capability appears in the CDB continuation segment (see 5.3).

The security parameters are described in 5.2.11.

If a QUERY command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the capacity quota.

6.26.2 Matches list parameter data format

The matches list parameter data (see table 126) contains the User_Object_ID for every user object that matched the query criteria.

Table 126 — Matches list parameter data format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
7	ADDITIONAL LENGTH (n-7)							_____ (LSB)
8	Reserved _____							
11	Reserved _____							
12	OBJECT DESCRIPTOR FORMAT (21h)						Reserved	
	Object descriptor list							
13	Object descriptor [first] (see table 111) _____							
	⋮							
n	Object descriptor [last] (see table 111) _____							

The ADDITIONAL LENGTH field indicates the number of bytes of matches list data that follow. If the matches list is truncated due to insufficient allocation length (see 6.26.1), the ADDITIONAL LENGTH field shall not be altered to reflect the truncation (i.e., the additional length indicates the number of bytes that would follow if the allocation length had

been infinite). If the untruncated number of bytes that follow is greater than FFFF FFFF FFFF FFFFh the additional length shall be set to FFFF FFFF FFFF FFFFh.

The OBJECT DESCRIPTOR FORMAT field shall contain 21h indicating that the object descriptors have the format shown in 6.20.3.3.

Each object descriptor (see 6.20.3.3) contains the User_Object_ID of one user object that matches the query criteria in the query list CDB continuation descriptor (see 5.4.3).

6.27 READ

The READ command (see table 127) requests that the device server return data to the application client from the specified user object.

Table 127 — READ command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (8885h)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID							
24	(MSB)							
31	USER_OBJECT_ID							
32	(MSB)							
39	LENGTH							
40	(MSB)							
47	STARTING BYTE ADDRESS							
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

The contents of the LENGTH field are described in 5.2.9. The data read from the user object shall be placed in the Data-In Buffer as described in 5.2.9.

The contents of the STARTING BYTE ADDRESS field are described in 5.2.12.

If the STARTING BYTE ADDRESS field specifies a byte that is beyond the user object logical length attribute value in the User Object Information attributes page (see 7.1.3.11), then:

- a) No bytes shall be transferred; and
- b) The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the CDB continuation segment (see 5.3), if any, contains a scatter/gather list CDB continuation descriptor and the STARTING BYTE ADDRESS field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the values in the LENGTH field and STARTING BYTE ADDRESS field result an attempt to read a byte that is beyond the user object logical length attribute value in the User Object Information attributes page, then:

- a) The bytes between the starting byte address and the user object logical length shall be transferred;
- b) The command shall be terminated with CHECK CONDITION status, with the sense key set to RECOVERED ERROR and the additional sense code set to READ PAST END OF USER OBJECT;
- c) The command-specific information sense data descriptor (see SPC-3) shall be included in the sense data; and
- d) The COMMAND-SPECIFIC INFORMATION field shall contain the number of bytes transferred.

Attempts to read bytes that have never been written shall result in zeros being returned.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field is not set to zero and the CDB continuation segment (see 5.3) contains a scatter/gather list CDB continuation descriptor, that descriptor shall be processed as described in 5.4.2.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if the CDB continuation segment (see 5.3):

- a) Contains more than one scatter/gather list CDB continuation descriptor; or
- b) Contains any CDB continuation descriptors other than the scatter/gather list CDB continuation descriptor.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.28 READ MAP

6.28.1 Introduction

The READ MAP command (see table 128) requests that the device server return a map of the data and attributes in the specified user object.

Table 128 — READ MAP command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (88B1h) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
14	(MSB) _____							
15	REQUESTED MAP TYPE _____ (LSB)							
16	(MSB) _____							
23	PARTITION_ID (see 5.2.10) _____ (LSB)							
24	(MSB) _____							
31	USER_OBJECT_ID (see 5.2.14) _____ (LSB)							
32	(MSB) _____							
39	ALLOCATION LENGTH (see 5.2.2) _____ (LSB)							
40	(MSB) _____							
47	DATA MAP BYTE OFFSET _____ (LSB)							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The REQUESTED MAP TYPE field (see table 129) specifies the map descriptor type values (see table 133) that shall be returned in the parameter data.

Table 129 — REQUESTED MAP TYPE field

Code	Description
0000h	Return all map type values.
0001h	Return only WRITTEN_DATA map type values.
0002h	Return only DATA_HOLE map type values.
0003h	Return only DAMAGED_DATA map type values.
0004h to 7FFFh	Reserved
8000h	Return only attributes map type values.
8001h to 8002h	Reserved
8003h	Return only DAMAGED_ATTRIBUTES map type values.
8004h to FFFFh	Reserved

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

The contents of the ALLOCATION LENGTH field are described in 5.2.2.

The DATA MAP BYTE OFFSET field specifies the first byte of user data to be represented in the returned map. If the DATA MAP BYTE OFFSET field specifies a byte that is beyond the user object logical length attribute value in the User Object Information attributes page (see 7.1.3.11), then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.28.2 READ MAP command and READ MAPS AND COMPARE command parameter data

The parameter data returned by a READ MAP command and a READ MAPS AND COMPARE command (see table 130) contains descriptors that describe the user data and attributes associated with a user object or objects.

Table 130 — READ MAP command and READ MAPS AND COMPARE command parameter data

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
7	ADDITIONAL LENGTH (n-7) _____ (LSB)							
8	Reserved _____							
15								
	Map descriptor list							
16	Map descriptor [first] _____							
23								
	⋮							
n-15	Map descriptor [last] _____							
n								

The ADDITIONAL LENGTH field indicates the number of bytes of READ MAP (see 6.28) command or READ MAPS AND COMPARE command (see 6.29) parameter data that follow. If the parameter data is truncated due to insufficient allocation length, the ADDITIONAL LENGTH field shall not be altered to reflect the truncation (i.e., the additional length indicates the number of bytes that would follow if the allocation length had been infinite). If the untruncated number of bytes that follow is greater than FFFF FFFF FFFF FFFFh the additional length shall be set to FFFF FFFF FFFF FFFFh.

Each map descriptor (see table 131) contains 16 bytes and provides information about user object attributes or one range of bytes within a user object's user data.

Table 131 — Map descriptor format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) MAP DESCRIPTOR INDEX (LSB)							
1								
2	(MSB) MAP DESCRIPTOR TYPE (LSB)							
3								
4	(MSB) DATA LENGTH (LSB)							
7								
8	(MSB) BYTE OFFSET (LSB)							
15								

In the parameter data for a READ MAP command the MAP DESCRIPTOR INDEX field is reserved. In the parameter data for the READ MAPS AND COMPARE command, the MAP DESCRIPTOR INDEX field indicates which user object the map descriptor represents as shown in table 132.

Table 132 — MAP DESCRIPTOR INDEX field

Value	Description
0000h	The user object specified by fields in the CDB
0001h	The user object specified by the first user object CDB continuation descriptor (see 5.4.4) in the CDB continuation segment (see 5.3)
0002h to FFFFh	Reserved

The MAP DESCRIPTOR TYPE field (see table 133) indicates the type of information this map descriptor contains.

Table 133 — MAP DESCRIPTOR TYPE field

Code	Name	Description
0000h		Reserved
0001h	WRITTEN_DATA	This map descriptor indicates the byte offset and data length of user data that has been written to stable storage (see 4.14) and is available for reading.
0002h	DATA_HOLE	This map descriptor indicates the byte offset and data length of a user data that lies between two WRITTEN_DATA regions, but for which no user data has been written.
0003h	DAMAGED_DATA	This map descriptor indicates the byte offset and data length of user data in which uncorrectable damage has been detected (see 4.11.3).
0004h	PAST_LAST_BYTE	This map descriptor is used by the READ MAPS AND COMPARE command (see 6.29) to indicate bytes that have other map descriptor types in other user objects but are beyond the user object logical length attribute value in the User Object Information attributes page (see 7.1.3.11) for the user object indicated by the MAP DESCRIPTOR INDEX field.
0005h to 8000h		Reserved
8001h	NORMAL_ATTRIBUTES	This map descriptor indicates that one or more user object attributes contain are undamaged. ^a
8002h		Reserved
8003h	DAMAGED_ATTRIBUTES	This map descriptor indicates that one or more user object attributes contain uncorrectable damage. ^a
8004h to FFFFh		Reserved
^a All the attributes in a user object are represented by a single map descriptor.		

If the map descriptor type is greater than 7FFFh, the BYTE OFFSET field is reserved. If the map descriptor type is less than 8000h, the BYTE OFFSET field indicates the starting byte address of the user data that this map descriptor represents. The byte offset in the first map descriptor shall be equal to or greater than the contents of the DATA MAP

BYTE OFFSET field in the CDB. The byte offset in any map descriptor after the first shall be greater than or equal to the byte offset in the preceding map descriptor.

If the map descriptor type is greater than 7FFFh, the DATA LENGTH field is reserved. If the map descriptor type is less than 8000h, the DATA LENGTH field indicates the number of bytes of user data, starting at byte offset, that this map descriptor represents.

The parameter data shall not contain any map descriptors in which the MAP DESCRIPTOR TYPE field is set to a value that is less than 8000h following the first map descriptor in which MAP DESCRIPTOR TYPE field is set to a value that is greater than 8000h.

6.29 READ MAPS AND COMPARE

The READ MAPS AND COMPARE command (see table 134) requests that the device server compare the map information that would be returned by a READ MAP command (see 6.28) for two user objects and return information about where the maps are different.

Table 134 — READ MAPS AND COMPARE command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (88B2h) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved	COMPARISON SCOPE		
12	TIMESTAMPS CONTROL							
13	Reserved _____							
15								
16	(MSB) _____							
23	PARTITION_ID (see 5.2.10) _____ (LSB)							
24	(MSB) _____							
31	USER_OBJECT_ID (see 5.2.14) _____ (LSB)							
32	(MSB) _____							
39	ALLOCATION LENGTH (see 5.2.2) _____ (LSB)							
40	(MSB) _____							
47	DATA MAP BYTE OFFSET _____ (LSB)							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	Get and set attributes parameters (see 5.2.6) _____							
79								
80	Capability (see 5.2.4) _____							
183								
184	Security parameters (see 5.2.11) _____							
235								

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The COMPARISON SCOPE field (see table 135) specifies the scope of the map comparison to be performed.

Table 135 — COMPARISON SCOPE field

Value	Description
000b	Only differences in the read map data shall be returned
001b	Differences in the read map data, and WRITTEN_DATA map descriptor types (see 6.28.2) where the data is not shared between the user objects shall be returned
010b to 111b	Reserved

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

The contents of the ALLOCATION LENGTH field are described in 5.2.2.

The DATA MAP BYTE OFFSET field specifies the first byte of user data to be compared. If the DATA MAP BYTE OFFSET field specifies a byte that is beyond the user object logical length attribute value in the User Object Information attributes page (see 7.1.3.11) of any user object involved in the comparison, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if the CDB continuation segment (see 5.3):

- a) Does not contain one extension capabilities CDB continuation descriptor (see 5.4.6);
- b) Does not contain one user object CDB continuation descriptor (see 5.4.4); or
- c) Contains any CDB continuation descriptors other than:
 - A) The extension capabilities CDB continuation descriptor; and
 - B) The user object CDB continuation descriptor.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4. The READ MAPS AND COMPARE command accesses two user objects. One capability is necessary for each user object accessed. One capability appears in the CDB. The other capability appears in the CDB continuation segment (see 5.3).

The security parameters are described in 5.2.11.

The parameter data returned by a READ MAPS AND COMPARE command is described in 6.28.2.

The map descriptors in the parameter data shall represent only bytes or attributes where the contents of the MAP DESCRIPTOR TYPE field (see 6.28.2) have different values among the user objects being compared under the requirements established by the COMPARISON SCOPE field.

The map descriptors in the parameter data with a map descriptor type (see 6.28.2) less than 8000h shall be sorted as follows:

- 1) Smallest to largest map descriptor index; and
- 2) Smallest to largest byte offset.

The map descriptors in the parameter data with a map descriptor type greater than 8000h shall be sorted from smallest to largest map descriptor index.

6.30 REFRESH SNAPSHOT OR CLONE

6.30.1 Introduction

The REFRESH SNAPSHOT OR CLONE command (see table 136) causes the OSD device server to:

- Restart the processing started by a CREATE SNAPSHOT command (see 6.10) or CREATE CLONE command (see 6.7) that was interrupted before completion; or
- Update the contents of a snapshot partition (see 4.13.2) to match the current contents of its source partition.

Table 136 — REFRESH SNAPSHOT OR CLONE command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (88ABh) _____ (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	IMMED_TR	Reserved	GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	FREEZE	Reserved			TIME OF DUPLICATION			
14	DUPLICATION METHOD							
15	Reserved							
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	Reserved _____							
47								
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	Get and set attributes parameters (see 5.2.6) _____							
79								
80	Capability (see 5.2.4) _____							
183								
184	Security parameters (see 5.2.11) _____							
235								

The REFRESH SNAPSHOT OR CLONE command accesses the following partitions:

- The destination partition that is specified by the PARTITION_ID field; and
- The source partition whose Partition_ID (see 4.6.4) is the value in the source partition attribute in the Snapshots Information attributes page (see 7.1.3.30) of the destination partition.

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The IMMED_TR bit is described in 5.2.5.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

If the FREEZE bit is set to zero, the REFRESH SNAPSHOT OR CLONE command shall not modify the contents of the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) of the source partition. If the FREEZE bit is set to one and source object freeze duplication management is supported (see 4.13.4.3), then the device server shall modify the contents of the object accessibility attribute in the Partition Information attributes page of the source partition as described in 6.30.2 and 6.30.4.

The TIME OF DUPLICATION field specifies which time of duplication source object management method (see 4.13.4.2) applies to the REFRESH SNAPSHOT OR CLONE command. If the TIME OF DUPLICATION field is set to DEFAULT (see table 44 in 4.13.4.2), then which time of duplication source object management method is used is specified as follows:

- a) If the partition type attribute in the Snapshots Information attributes page (see 7.1.3.30) of the destination partition is set to 01h (i.e., snapshot partition), then the default snapshot time of duplication method attribute in the Partition Information attributes page (see 7.1.3.9) of the source partition specifies which time of duplication management method applies to the REFRESH SNAPSHOT OR CLONE command; or
- b) If the partition type attribute in the Snapshots Information attributes page of the destination partition is set to 02h (i.e., clone partition), then the default clone time of duplication method attribute in the Partition Information attributes page of the source partition specifies which time of duplication management method applies to the REFRESH SNAPSHOT OR CLONE command.

The DUPLICATION METHOD field specifies which duplication method (see 4.13.3) applies to the REFRESH SNAPSHOT OR CLONE command. If the DUPLICATION METHOD field is set to DEFAULT (see table 43 in 4.13.3), then which duplication method is used is specified as follows:

- a) If the partition type attribute in the Snapshots Information attributes page (see 7.1.3.30) of the destination partition is set to 01h (i.e., snapshot partition), then the default snapshot duplication method attribute in the Partition Information attributes page (see 7.1.3.9) of the source partition specifies which duplication method applies to the REFRESH SNAPSHOT OR CLONE command; or
- b) If the partition type attribute in the Snapshots Information attributes page (see 7.1.3.30) of the destination partition is set to 02h (i.e., clone partition), then the default snapshot duplication method attribute in the Partition Information attributes page of the source partition specifies which duplication method applies to the REFRESH SNAPSHOT OR CLONE command.

The contents of PARTITION_ID field (i.e., the Partition_ID of the destination partition) are described in 5.2.10.

If the source partition attribute in the Snapshots Information attributes page (see 7.1.3.30) of the destination partition is undefined (see 3.1.51), the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if the CDB continuation segment (see 5.3):

- a) Does not contain one extension capabilities CDB continuation descriptor (see 5.4.6); or
- b) Contains any CDB continuation descriptors other than the extension capabilities CDB continuation descriptor.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4. The REFRESH SNAPSHOT OR CLONE command accesses two partitions. One capability is necessary for each partition accessed. One capability appears in the CDB. The other capability appears in the CDB continuation segment (see 5.3).

The security parameters are described in 5.2.11.

6.30.2 Processing before the IMMED_TR bit takes effect

A REFRESH SNAPSHOT OR CLONE command shall not be completed with GOOD status until at least all the operations described in this subclause have been performed. These operations shall be completed before completing the command with GOOD status even if the IMMED_TR bit is set to one.

If the snapshot forward attribute value in the Snapshots Information attributes page (see 7.1.3.30) of the destination partition is not equal to the source partition attribute value in the Snapshots Information attributes page of the destination partition and the support for refreshing attribute in the Root Information attributes page (see 7.1.3.8) contains MOST RECENT ONLY (i.e., 01h), then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the FREEZE bit is set to one and source object freeze duplication management (see 4.13.4.3) is not supported, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the requested time of duplication source object management method (see 4.13.4.2) is not supported or the requested duplication method (see 4.13.3) is not supported, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the partition type attribute in the Snapshots Information attributes page (see 7.1.3.30) of the destination partition contains 00h (i.e., primary partition), then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the of the destination partition contains 0000 0000h (i.e., allow all accesses), then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the active command status attribute is not set to zero in the Command Tracking attributes page (see 7.1.3.20) in snapshot/clone tracking well known collection (see 4.6.6.5.3) for the destination partition, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The refresh completion time in the Snapshots Information attributes page (see 7.1.3.30) of the destination partition shall be made undefined (see 3.1.51).

If the Snapshots Information attributes page (see 7.1.3.30) of the destination partition contains the following attribute values:

- a) The partition type attribute contains 01h (i.e., snapshot partition); and
- b) The snapshot forward attribute value is not equal to the source partition attribute value:

then:

- 1) The destination partition shall be unlinked from the history chain as described in 6.30.6; and
- 2) The destination partition shall be added as the newest entry in the history change as described in 6.30.5.

If the FREEZE bit is set to one, the device server shall:

- a) Note the value of the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the source partition for use in 6.30.4; and
- b) Set the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the source partition to 0000 0001h.

The object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the destination partition shall be set to 0000 0001h.

The snapshot/clone tracking well known collection (see 4.6.6.5.3) shall be update in the destination partition to include at least the following:

- a) Every user object and collection in the source partition shall have their User_Object_ID (see 4.6.5) or Collection_Object_ID (see 4.6.6) inserted as a member of the TRACKING collection (see 4.6.6.3); and
- b) The Command Tracking attributes page (see 7.1.3.20) shall be initialized to include at least the following:
 - A) The percent complete attribute shall be set to zero;
 - B) The active command status attribute shall be set to 88ABh (i.e., REFRESH SNAPSHOT OR CLONE command in progress); and
 - C) The ended command status attribute shall be set to FFFFh.

6.30.3 Processing after the IMMED_TR bit takes effect, if any

Every user object and collection in the source partition shall be duplicated in the destination snapshot partition using the:

- a) Duplication method (see 4.13.3) specified by the CDB; and
- b) Time of duplication method (see 4.13.4.2) specified by the CDB.

The membership and attributes of the snapshot/clone tracking well known collection for the destination partition should be maintained to restarting of an interrupted REFRESH SNAPSHOT OR CLONE command with the minimum of repeated work (e.g., user objects or collections that have been fully duplicated should be removed from the snapshot/clone tracking well known collection). Other factors (e.g., meeting the requirements of the END time of duplication method (see 4.13.4.2)) may cause user objects and collections to be added to the snapshot/clone tracking well known collection.

6.30.4 Command completion

When an error is encountered or when all user objects and collections in the source partition have been duplicated in the destination snapshot partition as described in 6.30.3, the REFRESH SNAPSHOT OR CLONE command processing shall be completed as described in this subclause.

If the FREEZE bit is set to one, the device server shall restore the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the source partition to the value noted in 6.30.2.

At least the following changes shall be made in the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the destination partition:

- a) The active command status attribute shall be set to zero;
- b) The ended command status attribute shall be set to indicate the condition (e.g., success or error) of the REFRESH SNAPSHOT OR CLONE command processing; and
- c) If sense data is available, it shall be placed in the sense data attribute.

If the REFRESH SNAPSHOT OR CLONE command processing complete (i.e., if the percent complete attribute in the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the destination partition is set to 100) and the ended command status attribute in the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the destination partition has been set to 0000h (i.e., GOOD status command completion), then:

- a) The refresh completion time attribute in the Snapshots Information attributes page (see 7.1.3.30) in the destination partition shall be set to the value of the clock attribute in the Root Information attributes page (see 7.1.3.8); and
- b) If the destination partition is a clone partition, then the object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the destination snapshot partition shall be set to 0000 0000h.

If the IMMED_TR bit is set to zero, status shall be returned for the REFRESH SNAPSHOT OR CLONE command.

6.30.5 Linking a snapshot as the most recent entry in the history chain

To add a snapshot partition in the history chain as the most recent snapshot, the snapshot backward and snapshot forward attributes in the Snapshots Information attributes page (see 7.1.3.30) of the following partitions shall be set as shown in table 137:

- a) Source partition (i.e., the partition whose Partition_ID (see 4.6.4) is the value in the source partition attribute of the snapshot partition to be added as the most recent snapshot);
- b) Destination partition (i.e., the snapshot partition to be added as the most recent snapshot); and
- c) Previous newest partition, if any (i.e., the partition whose Partition_ID in the snapshot backward attribute of the source partition before any changes are made).

Table 137 — Snapshot backward and forward attribute values set to add the most recent entry

Partition	Snapshots Information attributes page attribute	
	Snapshot backward	Snapshot forward
Source partition	The Partition_ID of the destination partition	not modified
Destination partition	The Partition_ID of the previous newest partition ^a	The Partition_ID of the source partition
Previous newest partition ^b	not modified	The Partition_ID of the destination partition
^a If the snapshot backward attribute of the source partition is undefined, this attribute shall also be undefined. ^b If the snapshot backward attribute of the source partition is undefined, the changes shown in this row are not made.		

6.30.6 Unlinking a snapshot from the history chain

To unlink a snapshot partition from the history chain, the snapshot backward and snapshot forward attributes in the Snapshots Information attributes page (see 7.1.3.30) of the following partitions shall be set as shown in table 138:

- a) Destination partition (i.e., the snapshot partition to be unlinked);
- b) Newer partition (i.e., the partition whose Partition_ID (see 4.6.4) is the value in the snapshot forward attribute of the snapshot partition to be unlinked before any changes are made); and
- c) Older partition, if any (i.e., the partition whose Partition_ID is the value in the snapshot backward attribute of the snapshot partition to be unlinked before any changes are made).

Table 138 — Snapshot backward and forward attribute values set to unlink an entry

Partition	Snapshots Information attributes page attribute	
	Snapshot backward	Snapshot forward
Newer partition	The value of the snapshot backward attribute of the destination partition ^a	not modified
Older partition	not modified	The value of the snapshot forward attribute of the destination partition ^a
^a If the specified attribute is undefined (see 3.1.51) the attribute to be set shall be made undefined.		

6.31 REMOVE

The REMOVE command (see table 139) removes a user object from any LINKED collections (see 4.6.6.2) in which it is a member, and deletes the user object.

Table 139 — REMOVE command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (888Ah) (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID (LSB)							
24	(MSB)							
31	USER_OBJECT_ID (LSB)							
32	Reserved							
47								
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5) (LSB)							
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.32 REMOVE COLLECTION

The REMOVE COLLECTION command (see table 140) removes a collection (see 4.6.6) from a partition.

Table 140 — REMOVE COLLECTION command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (8896h)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			FCR
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID							(LSB)
24	(MSB)							
31	COLLECTION_OBJECT_ID							(LSB)
32	Reserved							
47								
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							(LSB)
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The FCR (force collection removal) bit specifies the actions to be taken if the collection contains user objects. If the FCR bit is set to zero and the collection contains user objects, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to PARTITION OR COLLECTION CONTAINS USER OBJECTS. If the FCR bit is set to one, the collection shall be removed as follows even if it contains user objects:

- 1) The collection type (see 4.6.6.1) shall affect the processing of user objects in the as follows:
 - A) For LINKED collections (see 4.6.6.2), each user object in the collection shall be modified to indicate that the user object no longer is a member of the collection; or
 - B) For all other collection types, the user objects in the collection shall not be modified;
 and
- 2) The collection shall be removed.

If the FCR bit is set to one, the REMOVE COLLECTION command shall not be completed with GOOD status until all user object members of the collection have been modified as described in this subclause and the collection has been removed.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field (see 5.2.10) specify the Partition_ID of partition from which the collection is to be removed.

The contents of the COLLECTION_OBJECT_ID field specify the Collection_Object_ID (see 4.6.6) the collection to be removed.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB if the COLLECTION_OBJECT_ID field specifies:

- a) A well known collection (see 4.6.6.5); or
- b) A user tracking collection (see 3.1.55) in which the active command status attribute in the Command Tracking attributes page (see 7.1.3.20) is not set to zero.

Except for the processing of the FCR bit, the type of collection (see 4.6.6.1) being removed shall not affect how the REMOVE COLLECTION command is processed (e.g., the attributes of the Command Tracking attributes page (see 7.1.3.20), if any, shall not be modified).

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.33 REMOVE MEMBER OBJECTS

The REMOVE MEMBER OBJECTS command (see table 141) instructs the device server to remove all the user objects that are members of the specified user tracking collection (see 3.1.55). The REMOVE MEMBER OBJECTS command is a multi-object command (see 4.6.6.6).

Table 141 — REMOVE MEMBER OBJECTS command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (88A1h) _____ (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	IMMED_TR	Reserved	GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved _____							
15								
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	(MSB) _____							
31	COLLECTION_OBJECT_ID _____ (LSB)							
32	Reserved _____							
47								
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	Get and set attributes parameters (see 5.2.2) _____							
79								
80	Capability (see 5.2.4) _____							
183								
184	Security parameters (see 5.2.11) _____							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The IMMED_TR bit is described in 5.2.5.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The COLLECTION_OBJECT_ID field specifies Collection_Object_ID (see 4.6.6) to be processed. The device server shall constrain the Collection_Object_ID values as described in 4.6.6.6.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15. Get and set attributes processing requirements specific to multi-object commands are described in 4.6.6.6.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.34 REMOVE PARTITION

The REMOVE PARTITION command (see table 142) deletes a partition from the OSD logical unit.

Table 142 — REMOVE PARTITION command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (888Ch)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved	REMOVE SCOPE		
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID							
24	Reserved							
47								
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The REMOVE SCOPE field (see table 143) specifies the scope the partition removal operations requested with respect to user objects and collections within the partition.

Table 143 — REMOVE SCOPE field

Value	Description
000b	If there are any collections or user objects in the partition, the command shall be terminated with CHECK CONDITION status, the sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to PARTITION OR COLLECTION CONTAINS USER OBJECTS.
001b	If there are any collections or user objects in the partition, they shall be removed as part of removing the partition.
010b to 111b	Reserved

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB, if any of the following are true:

- a) The PARTITION_ID field is set to zero;
- b) The snapshots count attribute is defined (see 3.1.14) in the Snapshots Information attributes page (see 7.1.3.30) of the specified partition and contains a value other than zero; or
- c) The clones count attribute is defined (see 3.1.14) in the Snapshots Information attributes page of the specified partition and contains a value other than zero.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

GOOD status shall not be returned until:

- a) Any attempt to access the partition is assured to result in CHECK CONDITION status;
- b) An attempt to access any user object or collection, if any, in the partition is assured to result in CHECK CONDITION status; and
- c) A request to create a new partition with the same Partition_ID is assured not to fail due to a duplicate Partition_ID.

6.35 RESTORE PARTITION FROM SNAPSHOT

6.35.1 Introduction

The RESTORE PARTITION FROM SNAPSHOT command (see table 144) causes the OSD device server to update the contents of a main partition (i.e., primary partition or clone partition) to match the contents of a snapshot partition (see 4.13.2).

Table 144 — RESTORE PARTITION FROM SNAPSHOT command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (88ACh) _____ (LSB)							
10	Reserved			DPO	FUA	ISOLATION		
11	IMMED_TR	Reserved	GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
14	DUPLICATION METHOD							
15	Reserved							
16	(MSB) _____							
23	SNAPSHOT PARTITION_ID _____ (LSB)							
24	_____							
47	Reserved _____							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The RESTORE PARTITION FROM SNAPSHOT command accesses the following partitions:

- The snapshot partition that is specified by the SNAPSHOT PARTITION_ID field; and
- The main partition whose Partition_ID (see 4.6.4) is the value in the source partition attribute in the Snapshots Information attributes page (see 7.1.3.30) of the snapshot partition.

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The IMMEDIATE bit is described in 5.2.5.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The DUPLICATION METHOD field specifies which duplication method (see 4.13.3) applies to the RESTORE PARTITION FROM SNAPSHOT command. If the DUPLICATION METHOD field is set to DEFAULT (see table 43 in 4.13.3), then the default snapshot duplication method attribute in the Partition Information attributes page (see 7.1.3.9) of the main partition specifies which duplication method applies to the RESTORE PARTITION FROM SNAPSHOT command.

The SNAPSHOT PARTITION_ID field contains the Partition_ID (see 4.6.4) of the snapshot partition for the RESTORE PARTITION FROM SNAPSHOT command.

If the source partition attribute in the Snapshots Information attributes page (see 7.1.3.30) of the snapshot partition is undefined (see 3.1.51), the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if the CDB continuation segment (see 5.3):

- a) Does not contain one extension capabilities CDB continuation descriptor (see 5.4.6); or
- b) Contains any CDB continuation descriptors other than the extension capabilities CDB continuation descriptor.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4. The RESTORE PARTITION FROM SNAPSHOT command accesses two partitions. One capability is necessary for each partition accessed. One capability appears in the CDB. The other capability appears in the CDB continuation segment (see 5.3).

The security parameters are described in 5.2.11.

6.35.2 Processing before the IMMED_TR bit takes effect

A RESTORE PARTITION FROM SNAPSHOT command shall not be completed with GOOD status until at least all the operations described in this subclause have been performed. These operations shall be completed before completing the command with GOOD status even if the IMMED_TR bit is set to one.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB, if attributes in the Snapshots Information attributes page (see 7.1.3.30) of the snapshot partition have any of the following properties:

- a) The partition type attribute contains a value other than 01h (i.e., snapshot partition);
- b) The create completion time attribute is undefined (see 3.1.51) and the refresh completion time attribute is undefined.

If the requested duplication method (see 4.13.3) is not supported, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the main partition shall be set to 0000 0001h.

If the active command status attribute is not set to zero in the Command Tracking attributes page (see 7.1.3.20) in snapshot/clone tracking well known collection (see 4.6.6.5.3) for the snapshot partition, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The snapshot/clone tracking well known collection (see 4.6.6.5.3) shall be update in the snapshot partition to include at least the following:

- a) Every user object and collection in the snapshot partition shall have their User_Object_ID (see 4.6.5) or Collection_Object_ID (see 4.6.6) inserted as a member of the TRACKING collection (see 4.6.6.3); and
- b) The Command Tracking attributes page (see 7.1.3.20) shall be initialized to include at least the following:
 - A) The percent complete attribute shall be set to zero;
 - B) The active command status attribute shall be set to 88ACh (i.e., RESTORE PARTITION FROM SNAPSHOT command in progress); and
 - C) The ended command status attribute shall be set to FFFFh.

6.35.3 Processing after the IMMED_TR bit takes effect, if any

Every user object and collection in the snapshot partition shall be duplicated in the main snapshot partition using the duplication method (see 4.13.3) specified by the CDB.

The membership and attributes of the snapshot/clone tracking well known collection for the snapshot partition should be maintained to restarting of an interrupted RESTORE PARTITION FROM SNAPSHOT command with the minimum of repeated work (e.g., user objects or collections that have been fully duplicated should be removed from the snapshot/clone tracking well known collection).

6.35.4 Command completion

When an error is encountered or when all user objects and collections in the source partition have been duplicated in the destination snapshot partition as described in 6.35.3, the RESTORE PARTITION FROM SNAPSHOT command processing shall be completed as described in this subclause.

At least the following changes shall be made in the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the snapshot partition:

- a) The active command status attribute shall be set to zero;
- b) The ended command status attribute shall be set to indicate the condition (e.g., success or error) of the RESTORE PARTITION FROM SNAPSHOT command processing; and
- c) If sense data is available, it shall be placed in the sense data attribute.

If the RESTORE PARTITION FROM SNAPSHOT command processing complete (i.e., if the percent complete attribute in the Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the snapshot partition is set to 100) and the ended command status attribute in the

Command Tracking attributes page (see 7.1.3.20) of the snapshot/clone tracking well known collection (see 4.6.6.5.3) in the snapshot partition has been set to 0000h (i.e., GOOD status command completion), then:

- a) The restore completion time attribute in the Snapshots Information attributes page (see 7.1.3.30) in the main partition shall be set to the value of the clock attribute in the Root Information attributes page (see 7.1.3.8);
- b) The restore Partition_ID attribute in the Snapshots Information attributes page in the main partition shall be set to the Partition_ID of the snapshot partition;
- c) The object accessibility attribute in the Partition Information attributes page (see 7.1.3.9) in the main snapshot partition shall be set to 0000 0000h.

If the IMMED_TR bit is set to zero, status shall be returned for the RESTORE PARTITION FROM SNAPSHOT command.

6.36 SET ATTRIBUTES

The SET ATTRIBUTES command (see table 145) sets the specified attributes for the specified root object, partition, collection, or user object before attributes are retrieved (see 4.8.4).

Table 145 — SET ATTRIBUTES command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (888Fh)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID							
24	(MSB)							
31	USER_OBJECT_ID							
32	Reserved							
47								
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.37 SET KEY

The SET KEY command (see table 146) causes the OSD device server to update the specified secret key.

Table 146 — SET KEY command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8898h) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved		KEY TO SET	
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	Reserved				KEY VERSION			
25	(MSB) _____							
31	KEY IDENTIFIER _____ (LSB)							
32	(MSB) _____							
51	SEED _____ (LSB)							
52	_____							
79	Get and set attributes parameters (see 5.2.6) _____							
80	_____							
183	Capability (see 5.2.4) _____							
184	_____							
235	Security parameters (see 5.2.11) _____							

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The KEY TO SET field (see table 147) specifies which key shall be updated, which key identifier shall be stored, and which keys shall be invalid following the SET KEY command.

Table 147 — Key to set code values

Value	Key to update	Key identifier attribute to store	Keys to invalidate
00b	Reserved		
01b	Root	The root key identifier attribute in the Root Policy/Security attributes page (see 7.1.3.22)	Previous root key, and all partition and working keys
10b	Partition	The partition key identifier attribute in the Partition Policy/Security attributes page (see 7.1.3.23)	Previous partition key, and all working keys
11b	Working	The working key identifier attribute in the Partition Policy/Security attributes page selected by the KEY VERSION field in the CDB	None

For every key that is invalidated by a SET KEY command, the associated key identifier attribute shall have its attribute length set to zero.

The contents of the PARTITION_ID field are described in 5.2.10. If the KEY TO SET field contains 01b and the PARTITION_ID field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The KEY VERSION field specifies the working key version to be updated. If the KEY TO SET field contains 01b or 10b, the KEY VERSION field shall be ignored.

The KEY IDENTIFIER field specifies a unique identifier to be associated with the new key. Successful processing of the SET KEY command shall include storing the key identifier value in the attribute specified in table 147.

The SEED field contains a random number generated from a good source of entropy (e.g., as described in RFC 1750). The updated key values shall be computed as described in 4.12.9.2.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11. The secret key whose authentication key shall be used to compute the capability key for this SET KEY command is specified in 4.12.6.3.

6.38 SET MASTER KEY

6.38.1 Introduction

The SET MASTER KEY command (see table 148) causes the OSD device server to update the master key secret key.

Table 148 — SET MASTER KEY command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8899h) _____ (LSB)							
10	Reserved					ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved		DH_STEP	
12	TIMESTAMPS CONTROL							
13	Reserved _____							
23								
24	DH_GROUP							
25	(MSB) _____							
31	KEY IDENTIFIER _____ (LSB)							
32	(MSB) _____							
35	PARAMETER LIST LENGTH _____ (LSB)							
36	(MSB) _____							
39	ALLOCATION LENGTH _____ (LSB)							
40								
47	Reserved _____							
48	(MSB) _____							
51	CDB CONTINUATION LENGTH (see 5.2.5) _____ (LSB)							
52								
79	Get and set attributes parameters (see 5.2.6) _____							
80								
183	Capability (see 5.2.4) _____							
184								
235	Security parameters (see 5.2.11) _____							

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The DH_STEP (Diffie-Hellman step) field (see table 149) specifies which step in the Diffie-Hellman exchange to process.

Table 149 — Diffie-Hellman exchange step values

Value	Name	Reference
00b	SEED EXCHANGE	6.38.2
01b	CHANGE MASTER KEY	6.38.3
10b to 11b	Reserved	

If a SET MASTER KEY command is received with the DH_STEP field set to CHANGE MASTER KEY and no SET MASTER KEY command has been received with the DH_STEP field set to SEED EXCHANGE on the same I_T_L nexus during the past ten seconds, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

A device server that receives a SET MASTER KEY command on one I_T_L nexus while the processing the DH_steps on a different I_T_L nexus is incomplete may terminate the second SET MASTER KEY command with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to SYSTEM RESOURCE FAILURE.

The usage of other CDB fields is specified in the description of each Diffie-Hellman step.

6.38.2 Seed exchange

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The DH_GROUP field specifies the coded value selected from the Group Description list of coded values maintained by IANA (see <http://www.iana.org/assignments/ipsec-registry>) that identifies the DH_generator and DH_prime values to be used for the SEED EXCHANGE step. If the value in the DH_GROUP field does not appear in one of the DH group attributes in the Root Policy/Security attributes page (see 7.1.3.22) the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The KEY IDENTIFIER field is reserved for the SEED EXCHANGE step.

The PARAMETER LIST LENGTH field specifies the number of bytes of application client DH_data to be sent to the device server. The application client DH_data is computed as follows:

- 1) A random number, x , is generated having a value between 0 and DH_prime minus one observing the requirements in RFC 1750; and
- 2) The application client DH_data is equal to $\text{DH_generator}^x \text{ modulo DH_prime}$, where the DH_generator and DH_prime values are identified by the code value in the CDB DH_GROUP field.

The ALLOCATION LENGTH field specifies the number of bytes available to receive the device server DH_data (see table 150) sent in response to the SET MASTER KEY command. If the allocation length is not sufficient to contain

device server DH_data, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

Table 150 — Seed exchange device server DH_data format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
3	RESPONSE LENGTH (n-3)							(LSB)
4	DEVICE SERVER DH_DATA							
n								

The RESPONSE LENGTH field indicates the number of bytes of device server DH_data that follow.

The DEVICE SERVER DH_DATA field contains the DH_data computed by the device server as follows:

- 1) A random number, y , is generated having a value between 0 and DH_prime minus one observing the requirements in RFC 1750; and
- 2) The device server DH_data is equal to DH_generator^y modulo DH_prime, where the DH_generator and DH_prime values are identified by the code value in the CDB DH_GROUP field.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11. The master key authentication key shall be used to compute the capability key for this SET MASTER KEY command (see 4.12.6.3).

After GOOD status has been returned for the SET MASTER KEY command SEED EXCHANGE step and before the SET MASTER KEY command CHANGE MASTER KEY step is processed, the next authentication master key and next generation master key shall be computed as described in 4.12.9.2, using a seed value that is the concatenation of the following:

- 1) DH_generator^{xy} modulo DH_prime;
- 2) The contents of the OSD system ID attribute in the Root Information attributes page (see 7.1.3.8);
- 3) The contents of the product model attribute in the Root Information attributes page;
- 4) The contents of the serial number attribute in the Root Information attributes page;
- 5) The contents of the OSD name attribute in the Root Information attributes page; and
- 6) The contents of the username attribute in the Partition Information attributes page (see 7.1.3.9) for partition zero (see 3.1.33).

6.38.3 Change master key

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The DH_GROUP field is reserved for the CHANGE MASTER KEY step.

The KEY IDENTIFIER field specifies a unique identifier to be associated with the new master key. Successful processing of the SET MASTER KEY command CHANGE MASTER KEY step shall include storing the key identifier value in the master key identifier attribute in the Root Policy/Security attributes page (see 7.1.3.22).

The PARAMETER LIST LENGTH field specifies the number of bytes in the CHANGE MASTER KEY parameter list (see table 151). If the parameter list length causes truncation of any field in the CHANGE MASTER KEY parameter list, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to PARAMETER LIST LENGTH ERROR.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11. The next authentication master key computed after the return of GOOD status for the most recent SET MASTER KEY command SEED EXCHANGE step (see 6.38.2) shall be used to compute the capability key for this SET MASTER KEY command CHANGE MASTER KEY step (see 4.12.6.3).

Table 151 — Change master key DH_data format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
3	APPLICATION CLIENT DATA LENGTH (k-3)							(LSB)
4	APPLICATION CLIENT DH_DATA							
k								
k+1	(MSB)							
k+4	DEVICE SERVER DATA LENGTH (n-(k+4))							(LSB)
k+5	DEVICE SERVER DH_DATA							
n								

The APPLICATION CLIENT DATA LENGTH field specifies the number of bytes that follow in the APPLICATION CLIENT DH_DATA field.

The APPLICATION CLIENT DH_DATA field contains the application client DH_data from the SEED EXCHANGE step.

The DEVICE SERVER DATA LENGTH field specifies the number of bytes that follow in the DEVICE SERVER DH_DATA field.

The DEVICE SERVER DH_DATA field contains the device server DH_data from the SEED EXCHANGE step.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST if CHANGE MASTER KEY parameter data fails to compare in any of the following ways with the data exchanged in the SEED EXCHANGE step that was

most recently processed on this I_T_L nexus since a I_T nexus loss event, logical unit reset event, or reset event (see SAM-4), if any:

- a) The contents of the APPLICATION CLIENT DATA LENGTH field do not match the contents of the PARAMETER LIST LENGTH field in the SEED EXCHANGE step;
- b) The contents of the APPLICATION CLIENT DH_DATA field do not match the contents of the parameter data in the SEED EXCHANGE step;
- c) The contents of the DEVICE SERVER DATA LENGTH field do not match the contents of the RESPONSE LENGTH field in the SEED EXCHANGE step; or
- d) The contents of the DEVICE SERVER DH_DATA field do not match the contents of the DEVICE SERVER DH_DATA field in the SEED EXCHANGE step.

Successful processing of a SET MASTER KEY command CHANGE MASTER KEY step shall:

- a) Replace the authentication master key with the next authentication master key computed after the return of GOOD status for the most recent SET MASTER KEY command SEED EXCHANGE step (see 6.38.2);
- b) Replace the generation master key with the next generation master key computed after the return of GOOD status for the most recent SET MASTER KEY command SEED EXCHANGE step;
- c) Invalidate all of the following keys (see 4.12.9):
 - A) The root key;
 - B) The partition key for every partition on the OSD logical unit; and
 - C) Every working key in every partition on the OSD logical unit.

For every key that is invalidated by a SET MASTER KEY command CHANGE MASTER KEY step, the associated key identifier attribute shall have its attribute length set to zero.

6.39 SET MEMBER ATTRIBUTES

The SET MEMBER ATTRIBUTES command (see table 152) instructs the device server to set the specified attributes for the specified user tracking collection (see 3.1.55) and user object members of the user tracking collection before retrieving the attributes, if any, specified by the command (see 4.8.4). The SET MEMBER ATTRIBUTES command is a multi-object command (see 4.6.6.6).

Table 152 — SET MEMBER ATTRIBUTES command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (88A3h)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID							
24	(MSB)							
31	COLLECTION_OBJECT_ID							
32	Reserved							
47								
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
52	Get and set attributes parameters (see 5.2.2)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6. Page format attribute processing is illegal for the SET MEMBER ATTRIBUTES command. If the GET/SET CDBFMT field contains a value other than 11b, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The COLLECTION_OBJECT_ID field specifies Collection_Object_ID (see 4.6.6) to be processed. The device server shall constrain the Collection_Object_ID values as described in 4.6.6.6.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15. Get and set attributes processing requirements specific to multi-object commands are described in 4.6.6.6.

The same attributes (i.e., the same combinations of attribute page and attribute number) are retrieved and/or set for all user objects in the specified collection. If user object attributes are set, all such attributes are set to the same values in all user objects in the specified collection.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

6.40 WRITE

The WRITE command (see table 153) causes the specified number of bytes to be written to the specified user object at the specified relative location.

Table 153 — WRITE command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (8886h)							
10	Reserved			DPO	FUA	ISOLATION		
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
15								
16	(MSB)							
23	PARTITION_ID							
24	(MSB)							
31	USER_OBJECT_ID							
32	(MSB)							
39	LENGTH							
40	(MSB)							
47	STARTING BYTE ADDRESS							
48	(MSB)							
51	CDB CONTINUATION LENGTH (see 5.2.5)							
52	Get and set attributes parameters (see 5.2.6)							
79								
80	Capability (see 5.2.4)							
183								
184	Security parameters (see 5.2.11)							
235								

The contents of the DPO bit and the FUA bit are described in 5.2.3.

The contents of the ISOLATION field are described in 5.2.8.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.6.

The contents of the TIMESTAMPS CONTROL field are described in 5.2.13.

The contents of the PARTITION_ID field are described in 5.2.10.

The contents of the USER_OBJECT_ID field are described in 5.2.14.

The contents of the LENGTH field are described in 5.2.9. The data to be written to the user object shall be placed in the Data-Out Buffer as described in 5.2.9.

The contents of the STARTING BYTE ADDRESS field are described in 5.2.12. If the CDB continuation segment (see 5.3), if any, contains a scatter/gather list CDB continuation descriptor and the STARTING BYTE ADDRESS field contains a value other than zero, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

A WRITE to a byte that is greater than the value in the user object logical length attribute in the User Object Information attributes page (see 7.1.3.11) shall implicitly increase the value in the user object logical length attribute to the largest byte written.

The contents of the CDB CONTINUATION LENGTH field are described in 5.2.5. If the CDB CONTINUATION LENGTH field is not set to zero and the CDB continuation segment (see 5.3) contains a scatter/gather list CDB continuation descriptor, that descriptor shall be processed as described in 5.4.2.

The command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if the CDB continuation segment (see 5.3):

- a) Contains more than one scatter/gather list CDB continuation descriptor; or
- b) Contains any CDB continuation descriptors other than the scatter/gather list CDB continuation descriptor.

The get and set attributes parameters are described in 5.2.6. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.15.

The capability is described in 5.2.4.

The security parameters are described in 5.2.11.

If a WRITE command causes the value in the user object logical length attribute in the User Object Information attributes page (see 7.1.3.11) to exceed the value in the maximum user object length attribute in the User Object Quotas attributes page, then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the maximum user object length quota.

If a WRITE command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the capacity quota.

7 Parameters for OSD type devices

7.1 Attributes parameters

7.1.1 Attributes parameter formats

The following formats shall be provided for attributes parameter data:

- a) Page format (see 7.1.3); and
- b) List format (see 7.1.4).

Page format parameter data allows retrieval of attributes in formatted pages where only the attribute values appear in the parameter data.

Those attributes pages that do not have a defined page format are not accessible via page format parameter data (e.g., the Root Directory attributes page described in 7.1.3.4).

List format parameter data handles individual attributes in an identifier, length, value format, allowing access to any group of attributes in any order.

Attribute access is limited to the:

- a) Attributes associated with the OSD object addressed by the command; and
- b) Attributes in the Current Command attributes page (see 7.1.3.31).

NOTE 5 Addressing the root object allows access to the attributes associated with the root object and partition zero (see 3.1.33)

The format of the Data-In Buffer and Data-Out Buffer when attributes meta data is being used is described in 4.15.

A get attributes request for an attribute or attributes page having no previously established value shall not be considered an error. If an attribute value that has not been previously established is requested by specific attribute number, a list entry format value (see 7.1.4) with the ATTRIBUTE LENGTH field set to zero (see 4.8.2) shall be returned. If an attributes page that has no established definition is requested, a null attributes page (see 7.1.3.32) shall be returned.

7.1.2 Reporting illegal attempts to set attribute values

Attempts to set an attribute to an illegal value may be detected in the set attributes list or in certain CDB fields. Attempts to set an attribute to an illegal value shall be handled as follows:

- a) If the CDB ATTRIBUTE PAGE field, ATTRIBUTE NUMBER field, ATTRIBUTE LENGTH field, and ATTRIBUTE VALUE field (see 5.2.6.2) specify an illegal combination of attribute page, attribute number, attribute length, and attribute value, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB;
- b) If the CDB SET ATTRIBUTE PAGE field, SET ATTRIBUTE NUMBER field, and SET ATTRIBUTE LENGTH field (see 5.2.6.3) specify an illegal combination of attribute page, attribute number, and attribute length, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB;
- c) If the CDB SET ATTRIBUTE PAGE field, SET ATTRIBUTE NUMBER field, and SET ATTRIBUTE LENGTH field (see 5.2.6.3) specify an valid combination of attribute page, attribute number, and attribute length but the attribute value specified by the SET ATTRIBUTES OFFSET field is illegal, then the command shall be termi-

nated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST; or

- d) If a set attributes list (see 5.2.6.4) contains an entry that specifies an illegal combination of attribute page, attribute number, attribute length, and attribute value, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

7.1.3 OSD attributes pages

7.1.3.1 Attributes pages overview

Every attributes page is identified by a page number (see 4.8.5). Every attributes page includes attribute number 0h whose contents are described in 7.1.3.2.

In addition to attribute number 0h, an attributes page is composed of attribute values numbered 1h through FFFF FFEh.

The attributes pages defined by this standard are shown in table 154.

Table 154 — Attributes pages defined by this standard (part 1 of 2)

Page Number	Page Name	Page Format Defined	Support Requirements	Reference
0h	User Object Directory	No	Mandatory	7.1.3.7
1h	User Object Information	No	Mandatory	7.1.3.11
2h	User Object Quotas	Yes	Mandatory	7.1.3.14
3h	User Object Timestamps	Yes	Mandatory	7.1.3.18
4h	Collections	Yes	Optional ^a	7.1.3.21
5h	User Object Policy/Security	Yes	Mandatory	7.1.3.25
6h	User Object Error Recovery	Yes	Mandatory	7.1.3.29
7h to 7Fh	Reserved			
C+0h	Collection Directory	No	Optional ^a	7.1.3.6
C+1h	Collection Information	No	Optional ^a	7.1.3.10
C+2h	Reserved			
C+3h	Collection Timestamps	Yes	Optional ^a	7.1.3.17
C+4h	Command Tracking	No	Optional ^a	7.1.3.20
C+5h	Collection Policy/Security	Yes	Optional ^a	7.1.3.24
C+6h	Collection Error Recovery	Yes	Optional ^a	7.1.3.28
C+7h to C+7Fh	Reserved			
^a Support for this attributes page is mandatory if collections are supported (see 4.6.6).				
^b Support for this attributes page is mandatory if snapshots are supported (see 4.13.2).				

Table 154 — Attributes pages defined by this standard (part 2 of 2)

Page Number	Page Name	Page Format Defined	Support Requirements	Reference
P+0h	Partition Directory	No	Mandatory	7.1.3.5
P+1h	Partition Information	No	Mandatory	7.1.3.9
P+2h	Partition Quotas	Yes	Mandatory	7.1.3.13
P+3h	Partition Timestamps	Yes	Mandatory	7.1.3.16
P+4h	Attributes Access	No	Mandatory	7.1.3.19
P+5h	Partition Policy/Security	Yes	Mandatory	7.1.3.23
P+6h	Partition Error Recovery	Yes	Mandatory	7.1.3.27
P+7h	Snapshots Information	No	Optional ^b	7.1.3.30
P+8h to P+7Fh	Reserved			
R+0h	Root Directory	No	Mandatory	7.1.3.4
R+1h	Root Information	No	Mandatory	7.1.3.8
R+2h	Root Quotas	Yes	Mandatory	7.1.3.12
R+3h	Root Timestamps	Yes	Mandatory	7.1.3.15
R+4h	Reserved			
R+5h	Root Policy/Security	Yes	Mandatory	7.1.3.22
R+6h	Root Error Recovery	Yes	Mandatory	7.1.3.26
R+7h to R+7Fh	Reserved			
F000 0000h to FFFF FFFDh	Reserved			
FFFF FFFEh	Current Command	Yes	Mandatory	7.1.3.31
^a Support for this attributes page is mandatory if collections are supported (see 4.6.6).				
^b Support for this attributes page is mandatory if snapshots are supported (see 4.13.2).				

7.1.3.2 Attribute number 0h in all attributes pages

With the exception of the Root Directory, Partition Directory, and Collection Directory attributes pages, all attributes pages defined by this standard shall contain an identification of the page in attribute number 0h. All attributes pages should contain an identification of the page in attribute number 0h.

The format of the page identification shall be a 40 byte fixed length value with the format shown in table 155.

Table 155 — Attribute number 0h format for all attributes pages

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							(LSB)
7	VENDOR IDENTIFICATION							
8	(MSB)							(LSB)
39	ATTRIBUTES PAGE IDENTIFICATION							

The left-aligned, space-padded (see 3.8.2) VENDOR IDENTIFICATION field shall contain eight bytes of ASCII data (see 3.8) identifying the organization that has defined the contents of the attributes page. The format of the VENDOR

IDENTIFICATION field is identical to the format of the VENDOR IDENTIFICATION field in the standard INQUIRY data (see SPC-3).

NOTE 6 It is intended that the VENDOR IDENTIFICATION field provide a unique identification of the organization that defined the attributes page contents. In the absence of a formal registration procedure, T10 maintains a list of vendor identification codes in use (see SPC-3). Organizations are requested to voluntarily submit their identification codes to T10 to prevent duplication of codes. The T10 web site, www.t10.org, provides a convenient means to request an identification code.

The left-aligned, null-terminated, null-padded (see 3.8.2) ATTRIBUTES PAGE IDENTIFICATION field shall contain 32 bytes of ASCII data identifying the attributes page in which it appears.

If the VENDOR IDENTIFICATION field contains the ASCII characters "INCITS", the first characters in the ATTRIBUTES PAGE IDENTIFICATION field shall identify the INCITS technical committee that has defined the contents of the attributes page (e.g., attributes pages defined by this standard have the ASCII characters "T10" as the first characters in the ATTRIBUTES PAGE IDENTIFICATION field).

NOTE 7 Using the User Object Directory attributes page as an example, the VENDOR IDENTIFICATION field contains the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field contains the ASCII characters "T10 User Object Information". The attribute number 0h attribute value is "INCITS T10 User Object Directory".

7.1.3.3 Attribute number 0h for unidentified attributes pages

Certain attributes pages may be created dynamically making them subject to programming errors that fail to define an attribute number 0h for the attributes page as recommended by this standard. For such attributes pages, device servers use the unidentified page identification attribute value specified in this subclause.

The unidentified page identification attribute shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing eight ASCII space characters (i.e., 20h) and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "unidentified attributes page".

7.1.3.4 Root Directory attributes page

The Root Directory attributes page (R+0h) shall contain one attribute for every root attributes page number accessible via the logical unit.

Within the Root Directory attributes page:

- a) The attribute number of each attribute shall be equal to the page number of the accessible attributes page that it represents;
- b) The attribute value shall be equal to:
 - A) If the length of the attribute numbered 0h in the attributes page identified by the root directory attribute number is not zero, then the value of the attribute numbered 0h in the identified attributes page; or
 - B) If the length of the attribute numbered 0h in the attributes page identified by the root directory attribute number is zero, then:
 - a) If there are no attributes with a non-zero length in the attributes page identified by the root directory attribute number, then the root directory attribute value shall have a length of zero; or
 - b) If there are one or more attributes with a non-zero length in the attributes page identified by the root directory attribute number, then the root directory attribute shall have the unidentified page identification attribute value specified in 7.1.3.3.

The Root Directory page identification attribute (number R+0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Root Directory".

Attribute values in the Root Directory attributes page have the format described in 7.1.3.2.

Table 156 shows the attributes in the Root Directory attributes page when only the attributes pages defined in this standard are accessible via the logical unit.

Table 156 — Example Root Directory attributes page contents

Attribute Number	Attribute Value (ASCII characters)
R+0h	"INCITS T10 Root Directory"
R+1h	"INCITS T10 Root Information"
R+2h	"INCITS T10 Root Quotas"
R+3h	"INCITS T10 Root Timestamps"
R+5h	"INCITS T10 Root Policy/Security"
R+6h	"INCITS T10 Root Error Recovery"

The contents of the Root Directory attributes page shall be maintained by the OSD logical unit.

If a command attempts to set any attribute in the Root Directory attributes page, then the command shall be terminated as described in 7.1.2.

7.1.3.5 Partition Directory attributes page

The Partition Directory attributes page (P+0h) shall contain one attribute for every partition attributes page number accessible to the partition.

Within the Partition Directory attributes page:

- a) The attribute number of each attribute shall be equal to the page number of the accessible attributes page that it represents;
- b) The attribute value shall be equal to:
 - A) If the length of the attribute numbered 0h in the attributes page identified by the partition directory attribute number is not zero, then the value of the attribute numbered 0h in the identified attributes page; or
 - B) If the length of the attribute numbered 0h in the attributes page identified by the partition directory attribute number is zero, then:
 - a) If there are no attributes with a non-zero length in the attributes page identified by the partition directory attribute number, then the partition directory attribute value shall have a length of zero; or
 - b) If there are one or more attributes with a non-zero length in the attributes page identified by the partition directory attribute number, then the partition directory attribute shall have the unidentified page identification attribute value specified in 7.1.3.3.

The Partition Directory page identification attribute (number P+0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Partition Directory".

Attribute values in the Partition Directory attributes page have the format described in 7.1.3.2.

Table 157 shows the attributes in the Partition Directory attributes page when only the attributes pages defined in this standard are accessible via the logical unit.

Table 157 — Example Partition Directory attributes page contents

Attribute Number	Attribute Value (ASCII characters)
P+0h	"INCITS T10 Partition Directory"
P+1h	"INCITS T10 Partition Information"
P+2h	"INCITS T10 Partition Quotas"
P+3h	"INCITS T10 Partition Timestamps"
P+5h	"INCITS T10 Partition Policy/Security"
P+6h	"INCITS T10 Partition Error Recovery"

The contents of the Partition Directory attributes page shall be maintained by the OSD logical unit.

If a command attempts to set any attribute in the Partition Directory attributes page, then the command shall be terminated as described in 7.1.2.

7.1.3.6 Collection Directory attributes page

The Collection Directory attributes page (C+0h) shall contain one attribute for every collection attributes page number accessible to the collection.

Within the Collection Directory attributes page:

- a) The attribute number of each attribute shall be equal to the page number of the accessible attributes page that it represents;
- b) The attribute value shall be equal to:
 - A) If the length of the attribute numbered 0h in the attributes page identified by the collection directory attribute number is not zero, then the value of the attribute numbered 0h in the identified attributes page; or
 - B) If the length of the attribute numbered 0h in the attributes page identified by the collection directory attribute number is zero, then:
 - a) If there are no attributes with a non-zero length in the attributes page identified by the collection directory attribute number, then the collection directory attribute value shall have a length of zero; or
 - b) If there are one or more attributes with a non-zero length in the attributes page identified by the collection directory attribute number, then the collection directory attribute shall have the unidentified page identification attribute value specified in 7.1.3.3.

The Collection Directory page identification attribute (number C+0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Collection Directory".

Attribute values in the Collection Directory attributes page have the format described in 7.1.3.2.

Table 158 shows the attributes in the Collection Directory attributes page when only the attributes pages defined in this standard are accessible via the logical unit.

Table 158 — Example Collection Directory attributes page contents

Attribute Number	Attribute Value (ASCII characters)
C+0h	"INCITS T10 Collection Directory"
C+1h	"INCITS T10 Collection Information"
C+3h	"INCITS T10 Collection Timestamps"
C+4h	"INCITS T10 Command Tracking"
C+5h	"INCITS T10 Collection Policy/Security"
C+6h	"INCITS T10 Collection Error Recovery"

The contents of the Collection Directory attributes page shall be maintained by the OSD logical unit.

If a command attempts to set any attribute in the Collection Directory attributes page, then the command shall be terminated as described in 7.1.2.

7.1.3.7 User Object Directory attributes page

The User Object Directory attributes page (0h) shall contain one attribute for every user attributes page number accessible to the user object.

Within the User Object Directory attributes page:

- a) The attribute number of each attribute shall be equal to the page number of the accessible attributes page that it represents;
- b) The attribute value shall be equal to:
 - A) If the length of the attribute numbered 0h in the attributes page identified by the user object directory attribute number is not zero, then the value of the attribute numbered 0h in the identified attributes page; or
 - B) If the length of the attribute numbered 0h in the attributes page identified by the user object directory attribute number is zero, then:
 - a) If there are no attributes with a non-zero length in the attributes page identified by the user object directory attribute number, then the user object directory attribute value shall have a length of zero; or
 - b) If there are one or more attributes with a non-zero length in the attributes page identified by the user object directory attribute number, then the user object directory attribute shall have the unidentified page identification attribute value specified in 7.1.3.3.

The User Object Directory page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 User Object Directory".

Attribute values in the User Object Directory attributes page have the format described in 7.1.3.2.

Table 159 shows the attributes in the User Object Directory attributes page when only the attributes pages defined in this standard are accessible via the logical unit and collections are supported. If collections are not supported, attribute number 4h shall have a length of zero.

Table 159 — Example User Object Directory attributes page contents

Attribute Number	Attribute Value (ASCII characters)
0h	"INCITS T10 User Object Directory"
1h	"INCITS T10 User Object Information"
2h	"INCITS T10 User Object Quotas"
3h	"INCITS T10 User Object Timestamps"
4h	"INCITS T10 Collections"
5h	"INCITS T10 User Object Policy/Security"
6h	"INCITS T10 User Object Error Recovery"

The contents of the User Object Directory attributes page shall be maintained by the OSD logical unit.

If a command attempts to set any attribute in the User Object Directory attributes page, then the command shall be terminated as described in 7.1.2.

7.1.3.8 Root Information attributes page

The Root Information attributes page (R+1h) shall contain the attributes listed in table 160.

Table 160 — Root Information attributes page contents (part 1 of 2)

Attribute Number	Length (bytes) ^a	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h to 2h		Reserved	No	Yes
3h	20	OSD System ID	No	Yes
4h	8	Vendor identification	No	Yes
5h	16	Product identification	No	Yes
6h	32	Product model	No	Yes
7h	4	Product revision level	No	Yes
8h	variable	Product serial number	No	Yes
9h	variable	OSD name	Yes	No
Ah	8	Maximum CDB continuation length	No	Yes
Bh to 7Fh		Reserved	No	
80h	8	Total capacity	No	Yes
81h	8	Used capacity	No	Yes
82h		Reserved		
83h	4	Object accessibility	Yes	No
84h to BFh		Reserved	No	
C0h	8	Number of partitions	No	Yes
C1h to FFh		Reserved	No	
100h	6	Clock	No	Yes
101h to 10Fh		Reserved		
110h	1	Default isolation method	Yes	No
111h	32	Supported isolation methods	No	Yes
112h to 11Fh		Reserved		
120h	8	Data atomicity guarantee	No	Yes
121h	8	Data atomicity alignment	No	Yes
122h	8	Attributes atomicity guarantee	No	Yes
123h	1	Data/attributes atomicity multiplier	No	Yes
124h to 1C0h		Reserved	No	
^a A length of 0 in this column denotes an attribute that may be undefined (see 3.1.51).				

Table 160 — Root Information attributes page contents (part 2 of 2)

Attribute Number	Length (bytes) ^a	Attribute	Application Client Settable	OSD Logical Unit Provided
1C1h	0 or 4	Maximum snapshots count	No	Yes
1C2h	0 or 4	Maximum clones count	No	Yes
1C3h to 1CBh		Reserved	No	
1CCh	0 or 4	Maximum branch depth	No	Yes
1CDh to 1FFh		Reserved	No	
200h to 2FFh	0 or 4	Supported object duplication method	No	Yes
300h to 30Fh	0 or 4	Supported time of duplication method	No	Yes
310h	0 or 4	Support for duplicated object freezing	No	Yes
311h	0 or 1	Support for snapshot refreshing	No	Yes
312h to 0700 0000h		Reserved	No	
0700 0001h to 0700 FFFFh	0 or 4	Supported CDB continuation descriptor type	No	Yes
0701 0000h to FFFF FFFEh		Reserved	No	
^a A length of 0 in this column denotes an attribute that may be undefined (see 3.1.51).				

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Root Information".

The left-aligned, zero-padded (see 3.8.2) OSD system ID attribute (number 3h) shall contain an identification descriptor using the same format as defined for the Device Identification VPD page (see SPC-3) with the following additional requirements on the fields in the identification descriptor:

- a) The CODE SET field shall contain 1h (i.e., binary valued identifier);
- b) The PROTOCOL IDENTIFIER field shall contain Fh (i.e., not applicable to any specific protocol);
- c) The IDENTIFIER TYPE field shall contain one of the following values:
 - A) 1h (i.e., T10 vendor identification);
 - B) 2h (i.e., EUI-64 based); or
 - C) 3h (i.e., NAA);
- d) The ASSOCIATION field shall contain 0h (i.e., logical unit); and
- e) The identifier length shall be less than 17.

If the Device Identification VPD page contains an identification descriptor that meets the requirements for OSD System ID attribute value described in this subclause, the same identification descriptor should be used as the OSD System ID attribute value.

The vendor identification attribute (number 4h) shall contain the vendor identification of the manufacturer of the OBSD (see 3.1.27) in the same format as the VENDOR IDENTIFICATION field in the standard INQUIRY data (see SPC-3).

The product identification attribute (number 5h) shall contain the product identification of the OBSD in the same format as the PRODUCT IDENTIFICATION field in the standard INQUIRY data (see SPC-3).

The left-aligned, space-padded (see 3.8.2) product model attribute (number 6h) shall contain 32 bytes of ASCII characters (see 3.8.1) identifying the model of the OBSD.

The product revision level attribute (number 7h) shall contain the product revision level of the OBSD in the same format as the PRODUCT REVISION LEVEL field in the standard INQUIRY data (see SPC-3).

The product serial number attribute (number 8h) shall contain the product serial number of the OBSD in the same format as the PRODUCT SERIAL NUMBER field in the Unit Serial Number VPD page (see SPC-3).

The OSD name attribute (number 9h) shall contain an identification of the OSD logical unit specified by the application client. The OSD name attribute length shall be set to zero by a FORMAT OSD command (see 6.17).

The maximum CDB continuation length attribute (number Ah) shall contain the largest value allowed in the CDB CONTINUATION LENGTH field (see 5.2). If the device server does not support CDB continuation segments, then the maximum CDB continuation length attribute shall be set to zero. If the maximum CDB continuation length attribute contains a non-zero value, that value shall be at least 1 024.

The total capacity attribute (number 80h) shall contain the total number of bytes on the OSD logical unit.

The used capacity attribute (number 81h) shall contain the number of bytes used by all root object attributes, partitions, collections and user objects stored by the OSD logical unit including attributes bytes for the partition, collections, and user objects. If any objects in the OSD logical unit are the result of object duplications (see 4.13), the value of the used capacity attribute may increase for reasons that are not obvious consequences of the commands being processed as described in 4.13.5.

The object accessibility attribute (83h) specifies the accessibility of the root object, all partitions, all collections, and all user objects using one of the values shown in table 161. The object accessibility attribute shall be enforced as described in 4.7. The object accessibility attribute in the Root Information attributes page shall be set to zero (i.e., allow all accesses) by a FORMAT OSD command.

Table 161 — Object accessibility attribute values

Code	Description
0000 0000h	Allow all accesses
0000 0001h	Deny all write accesses and allow all read accesses
0000 0002h to FFFF FFFFh	Reserved

The number of partitions attribute (number C0h) shall contain the number of partitions present in the OSD logical unit.

The clock attribute (number 100h) shall contain the current time in use by the OSD device server represented as the count of the number of milliseconds elapsed since midnight, 1 January 1970 UT (see 3.1.52). The value shall be identical to the value of the adjustable clock attribute in the Root Policy/Security attributes page (see 7.1.3.22).

The default isolation method attribute (111h) specifies the default method for isolating the actions of one command from the actions of other concurrent commands using one of the values shown in table 162. The default isolation method attribute shall be set to one (i.e., NONE) by a FORMAT OSD command.

Table 162 — Default isolation method attribute values

Code	Name	Type	Description
00h			Reserved ^a
01h	NONE	M	The actions of one command are not isolated from the actions of other concurrent commands being processed by the device server.
02h	STRICT	M	The device server shall isolate all the actions of one command from the actions of other concurrent commands.
03h			Reserved
04h	RANGE	O	The device server shall isolate the actions of one command that modify a range of bytes within a user object from the actions of other concurrent commands that modify the same range of bytes within the same user object.
05h	FUNCTIONAL	O	The device server shall isolate the command function actions (see table 53 in 4.16.2.1) of one command from the same command function actions of other concurrent commands.
06h			Reserved
07h			Vendor specific
08h to FFh			Reserved ^a
Type Key: M = Command implementation is mandatory. O = Command implementation is optional.			
^a There is no way to represent codes in this range in the ISOLATION field (see 5.2.8).			

If a command attempts to set the default isolation method attribute to a value that the supported isolation methods attribute indicates is not supported, then the command shall be terminated as described in 7.1.2.

The supported isolation methods attribute (112h) is a bit mask (see table 163) that indicates which isolation methods (see table 162) are supported by the OBSD.

Table 163 — Supported isolation methods attribute contents

Bit Byte	7	6	5	4	3	2	1	0
0	VS	Reserved	FUNC	RANGE	Reserved	STRICT	NONE	Reserved
1	Reserved							
31								

If the FUNCTIONAL isolation method (see table 162) is supported, the FUNC bit shall be set to one. If the FUNCTIONAL isolation method is not supported, the FUNC bit shall be set to zero.

If the RANGE isolation method (see table 162) is supported, the RANGE bit shall be set to one. If the RANGE isolation method is not supported, the RANGE bit shall be set to zero.

The STRICT bit shall be set to one to indicate that the STRICT isolation method (see table 162) is supported.

The NONE bit shall be set to one to indicate that the NONE isolation method (see table 162) is supported.

The data atomicity guarantee attribute (120h), the data atomicity alignment attribute (121h), the attributes atomicity guarantee attribute (122h), and the data/attributes atomicity multiplier attribute (123h) are described in table 12 (see 4.9.2).

If it is defined (see 3.1.14), the maximum snapshots count attribute (number 1C1h) shall contain the non-zero number that is the largest value allowed in any snapshots count attribute in any Snapshots Information attributes page (see 7.1.3.30). If the maximum snapshots count attribute is defined, the following commands shall be supported:

- a) The CREATE SNAPSHOT command (see 6.10);
- b) The REFRESH SNAPSHOT command (see 6.30); and
- c) The RESTORE PARTITION FROM SNAPSHOT command (see 6.35).

If it is defined (see 3.1.14), the maximum clones count attribute (number 1C2h) shall contain the non-zero number that is the largest value allowed in any clones count attribute in any Snapshots Information attributes page (see 7.1.3.30). If the maximum clones count attribute is defined, the following commands shall be supported:

- a) The CREATE CLONE command (see 6.7); and
- b) The DETACH CLONE command (see 6.12).

If the CREATE SNAPSHOT command (see 6.10) is supported and the CREATE CLONE command (see 6.7) is supported, then the maximum branch depth attribute (number 1CCh) shall be defined (see 3.1.14) and shall contain largest value allowed in any branch depth attribute in any Snapshots Information attributes page (see 7.1.3.30).

Each supported object duplication method attribute (numbers 200h to 2FFh) shall contain support information for one object duplication method. The attribute number is 200h plus the code shown in table 43 (see 4.13.3) for the object duplication method for which support information is being provided. If an attribute in the range 200h to 2FFh is undefined (see 3.1.51), then the associated object duplication method is not supported for any usage. If an attribute in the range 200h to 2FFh is defined (see 3.1.14), then the attribute value (see table 164) indicates the functions for which the object duplication method is supported.

Table 164 — Supported object duplication method attributes contents

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							SNAPSHOT
1	Reserved							CLONE
2	Reserved							
3	Reserved							COPY_UO

If the SNAPSHOT bit is set to zero, the object duplication method indicated by the attribute number is not supported for use by the CREATE SNAPSHOT command (see 6.10). If the SNAPSHOT bit is set to one, the object duplication method indicated by the attribute number is supported for use by the CREATE SNAPSHOT command.

If the CLONE bit is set to zero, the object duplication method indicated by the attribute number is not supported for use by the CREATE CLONE command (see 6.7). If the CLONE bit is set to one, the object duplication method indicated by the attribute number is supported for use by the CREATE CLONE command.

If the COPY_UO bit is set to zero, the object duplication method indicated by the attribute number is not supported for use by the COPY USER OBJECTS command (see 6.4). If the COPY_UO bit is set to one, the object duplication method indicated by the attribute number is supported for use by the COPY USER OBJECTS command.

If any form of object duplication is supported (see 4.13), attribute number 200h (i.e., the supported object duplication method attribute for the DEFAULT object duplication method) and attribute number 2FFh (i.e., the supported object duplication method attribute for the DO NOT CARE object duplication method) shall be defined (see 3.1.14) and the attribute value shall be FFFF FFFFh (i.e., all uses of the DEFAULT object duplication method and the DO NOT CARE object duplication method shall be supported). The value of attribute number 200h or attribute number 2FFh should not be used to determine which object duplication commands are supported. This information is returned by the REPORT SUPPORTED OPERATION CODES command (see 6.23 and SPC-4).

Each supported time of duplication method attribute (numbers 300h to 30Fh) contains support information for one time of duplication source object management method. The attribute number is 300h plus the code shown in table 44 (see 4.13.4.2) for the time of duplication method for which support information is being provided. If an attribute in the range 300h to 30Fh is undefined (see 3.1.51), then the associated time of duplication method is not supported for any usage. If an attribute in the range 300h to 30Fh is defined (see 3.1.14), then the attribute value (see table 165) indicates the functions for which the object duplication method is supported.

Table 165 — Supported time of duplication method attributes contents

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							SNAPSHOT
1	Reserved							CLONE
2	Reserved							
3	Reserved							COPY_UO

If the SNAPSHOT bit is set to zero, the time of duplication method indicated by the attribute number is not supported for use by the CREATE SNAPSHOT command (see 6.10). If the SNAPSHOT bit is set to one, the time of duplication method indicated by the attribute number is supported for use by the CREATE SNAPSHOT command.

If the CLONE bit is set to zero, the time of duplication method indicated by the attribute number is not supported for use by the CREATE CLONE command (see 6.7). If the CLONE bit is set to one, the time of duplication method indicated by the attribute number is supported for use by the CREATE CLONE command.

If the COPY_UO bit is set to zero, the time of duplication method indicated by the attribute number is not supported for use by the COPY USER OBJECTS command (see 6.4). If the COPY_UO bit is set to one, the time of duplication method indicated by the attribute number is supported for use by the COPY USER OBJECTS command.

If any form of time of duplication source object management is supported (see 4.13.4.2), attribute number 300h (i.e., the supported time of duplication method attribute for the DEFAULT time of duplication method) and attribute number 308h (i.e., the supported time of duplication method attribute for the DO NOT CARE time of duplication method) shall be defined (see 3.1.14) and the attribute value shall be FFFF FFFFh (i.e., all uses of the DEFAULT time of duplication method and the DO NOT CARE time of duplication method) shall be supported if any time of duplication source object management is supported).

The support for duplicated object freezing attribute (number 310h) contains support information for source object freeze duplication management (see 4.13.4.3). If the support for duplicated object freezing attribute is undefined (see 3.1.51), then source object freeze duplication management is not supported for any usage. If the support for

duplicated object freezing attribute is defined (see 3.1.14), then the attribute value (see table 166) indicates the functions for which source object freeze duplication management is supported.

Table 166 — Support for duplicated object freezing attribute contents

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							SNAPSHOT
1	Reserved							CLONE
2	Reserved							
3	Reserved							COPY_UO

If the SNAPSHOT bit is set to zero, source object freeze duplication management (see 4.13.4.3) is not supported for use by the CREATE SNAPSHOT command (see 6.10). If the SNAPSHOT bit is set to one, source object freeze duplication management is supported for use by the CREATE SNAPSHOT command.

If the CLONE bit is set to zero, source object freeze duplication management (see 4.13.4.3) is not supported for use by the CREATE CLONE command (see 6.7). If the CLONE bit is set to one, source object freeze duplication management is supported for use by the CREATE CLONE command.

If the COPY_UO bit is set to zero, source object freeze duplication management (see 4.13.4.3) is not supported for use by the COPY USER OBJECTS command (see 6.4). If the COPY_UO bit is set to one, source object freeze duplication management is supported for use by the COPY USER OBJECTS command.

If it is defined (see 3.1.14), the support for snapshot refreshing attribute (number 311h) (see table 167) shall indicate how the REFRESH SNAPSHOT command (see 6.30) is supported. If the support for snapshot refreshing attribute is undefined (see 3.1.51), then the REFRESH SNAPSHOT command is not supported.

Table 167 — Support for snapshot refreshing attribute values

Value	Name	Description
00h	Reserved	
01h	MOST RECENT ONLY	The REFRESH SNAPSHOT command is allowed only if the value in the source partition attribute in the Snapshots Information attributes page (see 7.1.3.30) for the source partition is equal to the value in the snapshot forward attribute in the Snapshots Information attributes page.
02h to FEh	Reserved	
FFh	UNLIMITED	The REFRESH SNAPSHOT command has no limits on the source partition.

Each supported CDB continuation descriptor type attribute (numbers 0700 0001h to 0700 FFFFh) contains support information for one CDB continuation descriptor type (see 5.4.1). The attribute number is 0700 0000h plus the CDB continuation descriptor type value shown in table 71 (see 5.4.1). If the supported CDB continuation descriptor type attribute is undefined (see 3.1.51), then the CDB continuation descriptor type is not supported for any usage. If the supported CDB continuation descriptor type attribute is defined (see 3.1.14), then the attribute value (see table 168) indicates the level of support for the CDB continuation descriptor type.

Table 168 — Supported CDB continuation descriptor type attribute contents

Value	Description
0000 0000h	The CDB continuation descriptor type is not supported for any usage.
0000 0001h to 0000 0007h	Reserved
0000 0008h to FFFF FFD0h	The CDB continuation descriptor type is supported only if the contents of the CDB CONTINUATION DESCRIPTOR LENGTH field are less than or equal to the value of the supported CDB continuation descriptor type attribute.
FFFF FFD1h to FFFF FFFEh	Reserved
FFFF FFFFh	The CDB continuation descriptor type is supported for all forms of usage.

If a command attempts to set an attribute that table 160 states is not application client settable, then the command shall be terminated as described in 7.1.2.

7.1.3.9 Partition Information attributes page

The Partition Information attributes page (P+1h) shall contain the attributes listed in table 169.

Table 169 — Partition Information attributes page contents

Attribute Number	Length (bytes) ^a	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	8	Partition_ID	No	Yes
2h to 8h		Reserved	No	
9h	variable	Username	Yes	No
Ah to 80h		Reserved	No	
81h	8	Used capacity	No	Yes
82h		Reserved		
83h	4	Object accessibility	Yes	No
84h	0 or 8	Potential used capacity increment	No	Yes
85h to BFh		Reserved	No	
C1h	8	Number of collections and user objects	No	Yes
C2h to D0h		Reserved	No	
D1h	0 or 8	Actual data space	No	Yes
D2h	0 or 8	Reserved data space	Yes	No
D3h to 1FFh		Reserved	No	
200h	0 or 4	Default snapshot duplication method	Yes	No
201h	0 or 4	Default clone duplication method	Yes	No
202h	0 or 4	Default copy user objects duplication method	Yes	No
203 to 1FFh		Reserved	No	
300h	0 or 4	Default snapshot time of duplication method	Yes	No
301h	0 or 4	Default clone time of duplication method	Yes	No
302h	0 or 4	Default copy user objects time of duplication method	Yes	No
303h to FFFF FFEh		Reserved	No	
^a A length of 0 in this column denotes an attribute that may be undefined (see 3.1.51).				

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Partition Information".

The Partition_ID attribute (number 1h) shall contain the Partition_ID of the partition with which the Partition Information attributes page is associated.

The username attribute (number 9h) shall contain a binary valued identification of the user of the partition specified by the application client. A CREATE PARTITION command (see 6.8) shall copy the username attribute from the Partition Information attributes page for partition zero (see 3.1.33) to the new Partition Information attributes page.

For all partitions except partition zero, the used capacity attribute (number 81h) shall contain the number of allocated bytes for the partition as described in this subclause. For partition zero, the used capacity attribute shall contain the number of allocated bytes for partition zero and all other partitions described in this subclause. The number of allocated bytes shall be computed as the sum of the following:

- a) The number of bytes used by:
 - A) The partition or partitions;
 - B) All collections within the partition or partitions;
 - C) All user tracking collections within the partition or partitions; and
 - D) All user objects within the partition or partitions including attributes bytes;
 and
- b) The number of unused reserved bytes computed as:
 - A) Value in the reserved data space attribute in this Partition Information attributes page minus the value in the actual data space attribute in this Partition Information attributes page; or
 - B) Zero if the value in the actual data space attribute in this Partition Information attributes page is larger than the value in the reserved data space attribute in this Partition Information attributes page.

If any object in the partition the result of object duplications (see 4.13), the value of the used capacity attribute may increase for reasons that are not obvious consequences of the commands being processed as described in 4.13.5.

The object accessibility attribute (83h) specifies the accessibility of the partition, all collections in the partition, and all user objects in the partition using one of the values shown in table 161 (see 7.1.3.8). The object accessibility attribute shall be enforced as described in 4.7. The object accessibility attribute in the Partition Information attributes page shall be set to zero (i.e., allow all accesses) by a CREATE PARTITION command.

The potential used capacity increment (number 84h) shall contain the maximum number of bytes by which the used capacity attribute in this Partition Information attributes page might increase due to ongoing command processing as described in 4.13.5.

For all partitions except partition zero, the number of collections and user objects attribute (number C1h) shall contain the sum of the number of collections and the number of user objects in the partition. For partition zero, the number of collections and user objects attribute shall contain zero.

If the OSD logical unit does not support the reserved data space attribute, the actual data space attribute (D1h) shall be undefined (see 3.1.51). If the reserved data space attribute is supported, the actual data space attribute shall be defined (see 3.1.14) and contain the number of bytes used by all user objects in the partition to store data transferred in the command data or parameter data segment of the Data-Out Buffer (see 4.15.4) by APPEND commands (see 6.2), CLEAR commands (see 6.3), CREATE AND WRITE commands (see 6.6), and/or WRITE commands (see 6.40) to a user object.

If the reserved data space attribute (D2h) is defined (see 3.1.14), it contains the minimum value that the actual data space attribute in this Partition Information attributes page is allowed to reach before an APPEND command, a CLEAR command, a CREATE AND WRITE command, or a WRITE command may be terminated with the ADDITIONAL SENSE CODE field set to 55h (i.e., SYSTEM RESOURCE FAILURE).

If snapshot object duplication is supported (see 4.13), the default snapshot duplication method attribute (number 200h) shall be defined (see 3.1.14) and shall contain one of the codes in table 43 (see 4.13.3) other than DEFAULT.

A CREATE PARTITION command (see 6.9) shall set the default snapshot duplication method attribute to DO NOT CARE (see table 43). If a command attempts to set the default snapshot duplication method attribute to DEFAULT or to a code that the supported object duplication method attributes in the Root Information attributes page (see 7.1.3.8) indicate is not supported, then the command shall be terminated as described in 7.1.2, and the value in the default snapshot duplication method attribute shall not be changed.

If clone object duplication is supported (see 4.13), the default clone duplication method attribute (number 201h) shall be defined (see 3.1.14) and shall contain one of the codes in table 43 (see 4.13.3) other than DEFAULT. A CREATE PARTITION command (see 6.9) shall set the default clone duplication method attribute to DO NOT CARE (see table 43). If a command attempts to set the default clone duplication method attribute to DEFAULT or to a code that the supported object duplication method attributes in the Root Information attributes page (see 7.1.3.8) indicate is not supported, then the command shall be terminated as described in 7.1.2, and the value in the default clone duplication method attribute shall not be changed.

If the copy user objects form of object duplication is supported (see 4.13), the default copy user objects duplication method attribute (number 202h) shall be defined (see 3.1.14) and shall contain one of the codes in table 43 (see 4.13.3) other than DEFAULT. A CREATE PARTITION command (see 6.9) shall set the default copy user objects duplication method attribute to DO NOT CARE (see table 43). If a command attempts to set the default copy user objects duplication method attribute to DEFAULT or to a code that the supported object duplication method attributes in the Root Information attributes page (see 7.1.3.8) indicate is not supported, then the command shall be terminated as described in 7.1.2, and the value in the default copy user objects duplication method attribute shall not be changed.

If snapshot object duplication is supported (see 4.13), the default snapshot time of duplication method attribute (number 300h) shall be defined (see 3.1.14) and shall contain one of the codes in table 44 (see 4.13.4.2) other than DEFAULT. A CREATE PARTITION command (see 6.9) shall set the default snapshot time of duplication method attribute to DO NOT CARE (see table 44). If a command attempts to set the default snapshot time of duplication method attribute to DEFAULT or to a code that the supported object duplication method attributes in the Root Information attributes page (see 7.1.3.8) indicate is not supported, then the command shall be terminated as described in 7.1.2, and the value in the default snapshot time of duplication method attribute shall not be changed.

If clone object duplication is supported (see 4.13), the default clone time of duplication method attribute (number 301h) shall be defined (see 3.1.14) and shall contain one of the codes in table 44 (see 4.13.4.2) other than DEFAULT. A CREATE PARTITION command (see 6.9) shall set the default clone time of duplication method attribute to DO NOT CARE (see table 44). If a command attempts to set the default clone time of duplication method attribute to DEFAULT or to a code that the supported object duplication method attributes in the Root Information attributes page (see 7.1.3.8) indicate is not supported, then the command shall be terminated as described in 7.1.2, and the value in the default clone time of duplication method attribute shall not be changed.

If the copy user objects form of object duplication is supported (see 4.13), the default copy user objects time of duplication method attribute (number 303h) shall be defined (see 3.1.14) and shall contain one of the codes in table 44 (see 4.13.4.2) other than DEFAULT. A CREATE PARTITION command (see 6.9) shall set the default copy user objects time of duplication method attribute to DO NOT CARE (see table 44). If a command attempts to set the default copy user objects time of duplication method attribute to DEFAULT or to a code that the supported object duplication method attributes in the Root Information attributes page (see 7.1.3.8) indicate is not supported, then the command shall be terminated as described in 7.1.2, and the value in the default copy user objects time of duplication method attribute shall not be changed.

If a command attempts to set an attribute that table 169 states is not application client settable, then the command shall be terminated as described in 7.1.2.

7.1.3.10 Collection Information attributes page

The Collection Information attributes page (C+1h) shall contain the attributes listed in table 170.

Table 170 — Collection Information attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	8	Partition_ID	No	Yes
2h	8	Collection_Object_ID	No	Yes
3h to 8h		Reserved	No	
9h	variable	Username	Yes	No
Ah	1	Collection type	No	Yes
Bh to 80h		Reserved	No	
81h	8	Used capacity	No	Yes
82h		Reserved		
83h	4	Object accessibility	Yes	No
84h to FFFF FFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Collection Information".

The Partition_ID attribute (number 1h) shall contain the Partition_ID of the collection with which the Collection Information attributes page is associated.

The Collection_Object_ID attribute (number 2h) shall contain the Collection_Object_ID (see 4.6.6) of the collection with which the Collection Information attributes page is associated.

The username attribute (number 9h) shall contain a binary valued identification of the user for the collection specified by the application client. A CREATE COLLECTION command (see 6.8) shall copy the username attribute from the Partition Information attributes page (see 7.1.3.9) to the new Collection Information attributes page.

The collection type attribute (number Ah) shall identify the characteristics (see table 171) of the collection.

Table 171 — Collection type codes

Code	Name	Description
00h	LINKED	User objects may be added to or removed from the collection using the Collections attributes page (see 7.1.3.21).
01h	TRACKING	Changes in the Collections attributes page (see 7.1.3.21) shall not affect the membership of TRACKING type collections. Changes in the membership of TRACKING type collections shall not affect the attributes in the Collections attributes page. The membership of TRACKING type collections is maintained by the device server based on processing requested by the application client (e.g., the processing of multi-object commands (see 4.6.6.6)).
02h to Eeh		Reserved
EFh	SPONTANEOUS	The membership of a SPONTANEOUS type collection shall be recomputed every time the collection is accessed. All SPONTANEOUS collections are well known collections (see 4.6.6.5) and the collection's Collection_Object_ID specifies how to compute the collection's membership.
F0h to FFh		Reserved

The used capacity attribute (number 81h) shall contain the number of bytes used by the collection including attributes bytes. If the collection the result of an object duplication (see 4.13), the value of the used capacity attribute may increase for reasons that are not obvious consequences of the commands being processed as described in 4.13.5.

The object accessibility attribute (83h) specifies the accessibility of the collection using one of the values shown in table 161 (see 7.1.3.8). The object accessibility attribute shall be enforced as described in 4.7. The object accessibility attribute in the Collection Information attributes page shall be set to zero (i.e., allow all accesses) by a CREATE COLLECTION command.

If a command attempts to set an attribute that table 170 states is not application client settable, then the command shall be terminated as described in 7.1.2.

7.1.3.11 User Object Information attributes page

The User Object Information attributes page (1h) shall contain the attributes listed in table 172.

Table 172 — User Object Information attributes page contents

Attribute Number	Length (bytes) ^a	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	8	Partition_ID	No	Yes
2h	8	User_Object_ID	No	Yes
3h to 8h		Reserved	No	
9h	variable	Username	Yes	No
Ah to 80h		Reserved	No	
81h	8	Used capacity	No	Yes
82h	8	User object logical length	Yes	Yes
83h	4	Object accessibility	Yes	No
84h to D0h		Reserved	No	
D1h	0 or 8	Actual data space	No	Yes
D2h	0 or 8	Reserved data space	Yes	No
D3h to FFFF FFEh		Reserved	No	
^a A length of 0 in this column denotes an attribute that may be undefined (see 3.1.51).				

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 User Object Information".

The Partition_ID attribute (number 1h) shall contain the Partition_ID of the user object with which the User Object Information attributes page is associated.

The User_Object_ID attribute (number 2h) shall contain the User_Object_ID of the user object with which the User Object Information attributes page is associated.

The username attribute (number 9h) shall contain a binary valued identification of the user for the user object specified by the application client. A CREATE command (see 6.5) or CREATE AND WRITE command (see 6.6) shall copy the username attribute from the Partition Information attributes page (see 7.1.3.9) to the new User Object Information attributes page.

The used capacity attribute (number 81h) shall contain the sum of:

- a) The number of bytes used by the user object including attributes bytes; and
- b) The number of unused reserved bytes computed as:
 - A) Value in the reserved data space attribute in this User Object Information attributes page minus the value in the actual data space attribute in this User Object Information attributes page; or

- B) Zero if the value in the actual data space attribute in this User Object Information attributes page is larger than the value in the reserved data space attribute in this User Object Information attributes page.

If the user object the result of an object duplication (see 4.13), the value of the used capacity attribute may increase for reasons that are not obvious consequences of the commands being processed as described in 4.13.5.

The user object logical length attribute (number 82h) specifies the largest valued byte number written in the associated user object. Setting the user object logical length attribute to a value that is smaller than the user object's logical length known to the OSD device server shall cause the user object to be truncated to the specified length. Setting the user object logical length attribute to a value that is larger than the user object's logical length known to the OSD device server shall cause unwritten bytes to be added at the end of the user object.

At attempt to set the user object logical length attribute a value that is larger than the value in the maximum user object length attribute in the User Object Quotas attributes page (see 7.1.3.14) shall generate a quota error (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the maximum user object length quota.

If setting the user object logical length attribute to a value that is larger than the user object's logical length known to the OSD device server causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the capacity quota.

The object accessibility attribute (83h) specifies the accessibility of the user object using one of the values shown in table 161 (see 7.1.3.8). The object accessibility attribute shall be enforced as described in 4.7. The object accessibility attribute in the User Object Information attributes page shall be set to zero (i.e., allow all accesses) by a CREATE command or a CREATE AND WRITE command.

If the OSD logical unit does not support the reserved data space attribute, the actual data space attribute (D1h) shall be undefined (see 3.1.51). If the reserved data space attribute is supported, the actual data space attribute shall be defined (see 3.1.14) and shall contain the number of bytes used by the user object to store data transferred in the command data segment of the Data-Out Buffer (see 4.15.4) by APPEND commands (see 6.2), CLEAR commands (see 6.3), CREATE AND WRITE commands (see 6.6), and/or WRITE commands (see 6.40) to the user object.

If the reserved data space attribute (D2h) is defined (see 3.1.14), it contains the minimum value that the actual data space attribute in this User Object Information attributes page is allowed to reach before an APPEND command, a CLEAR command, a CREATE AND WRITE command, or a WRITE command may be terminated with the ADDITIONAL SENSE CODE field set to 55h (i.e., SYSTEM RESOURCE FAILURE).

If a command attempts to set an attribute that table 172 states is not application client settable, then the command shall be terminated as described in 7.1.2.

7.1.3.12 Root Quotas attributes page

The Root Quotas attributes page (R+2h) shall contain the attributes listed in table 173. Except for attribute number 0h, all attributes in the Root Quotas attributes page are quotas (see 4.10).

Table 173 — Root Quotas attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	8	Default maximum user object length	Yes	No
2h to 1 0000h		Reserved	No	
1 0001h	8	Partition capacity quota	Yes	No
1 0002h	8	Partition object count	Yes	No
1 0003h to 1 0080h		Reserved	No	
1 0081h	4	Partition collections per user object	Yes	No
1 0082h to 2 0001h		Reserved	No	
2 0002h	8	Partition count	Yes	No
2 0003h to FFFF FFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Root Quotas".

The default maximum user object length attribute (number 1h) specifies the value to be copied to the default maximum user object length attribute in the Partition Quotas attributes page (see 7.1.3.13) for each partition, when it is created. The FORMAT OSD command (see 6.17) shall set the default maximum user object length attribute to FFFF FFFF FFFF FFFFh.

The partition capacity quota attribute (number 1 0001h) specifies the value to be copied to the capacity quota attribute in the Partition Quotas attributes page for each partition, when it is created. The FORMAT OSD command shall set the partition capacity quota attribute to FFFF FFFF FFFF FFFFh.

The partition object count attribute (number 1 0002h) specifies the value to be copied to the object count attribute in the Partition Quotas attributes page for each partition, when it is created. The FORMAT OSD command shall set the partition object count attribute to FFFF FFFF FFFF FFFFh.

The partition collections per user object attribute (number 1 0081h) specifies the value to be copied to the collections per user object attribute in the Partition Quotas attributes page for each partition, when it is created. The FORMAT OSD command shall set the partition collections per user object attribute to FFFF FFFFh.

The partition count attribute (number 2 0002h) specifies the maximum value allowed in the number of partitions attribute in the Root Information attributes page (see 7.1.3.8). If a CREATE PARTITION command (see 6.9) attempts to exceed the partition count quota, a quota error (see 4.10.2) shall be generated. The FORMAT OSD command shall set the partition count attribute to FFFF FFFF FFFF FFFFh. If a command attempts to set the partition count attribute value to zero, then the command shall be terminated as described in 7.1.2, and the value in the partition count attribute shall not be changed.

If a command attempts to set an attribute that table 173 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Root Quotas attributes page is shown in table 174.

Table 174 — Root Quotas attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	PAGE NUMBER (R+2h)						(LSB)
3								
4	(MSB)	PAGE LENGTH (24h)						(LSB)
7								
8	(MSB)	DEFAULT MAXIMUM USER OBJECT LENGTH						(LSB)
15								
16	(MSB)	PARTITION CAPACITY QUOTA						(LSB)
23								
24	(MSB)	PARTITION OBJECT COUNT						(LSB)
31								
32	(MSB)	PARTITION COLLECTIONS PER USER OBJECT						(LSB)
35								
36	(MSB)	PARTITION COUNT						(LSB)
43								

The PAGE NUMBER field contains the attributes page number of the Root Quotas attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Root Quotas attributes page.

The DEFAULT MAXIMUM USER OBJECT LENGTH field contains the value of the default maximum user object length attribute.

The PARTITION CAPACITY QUOTA field contains the value of the partition capacity quota attribute.

The PARTITION OBJECT COUNT field contains the value of the partition object count attribute.

The PARTITION COLLECTIONS PER USER OBJECT field contains the value of the partition collections per user object attribute.

The PARTITION COUNT field contains the value of the partition count attribute.

7.1.3.13 Partition Quotas attributes page

The Partition Quotas attributes page (P+2h) shall contain the attributes listed in table 175. Except for attribute number 0h, all attributes in the Partition Quotas attributes page are quotas (see 4.10).

Table 175 — Partition Quotas attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	8	Default maximum user object length	Yes	No
2h to 1 0000h		Reserved	No	
1 0001h	8	Capacity quota	Yes	No
1 0002h	8	Object count	Yes	No
1 0003h to 1 0080h		Reserved	No	
1 0081h	4	Collections per user object	Yes	No
1 0082h to FFFF FFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Partition Quotas".

The default maximum user object length attribute (number 1h) specifies the value to be copied to the maximum user object length attribute in the User Object Quotas attributes page (see 7.1.3.14) for each user object, when it is created. The CREATE PARTITION command (see 6.9) shall set this attribute to the value in the default maximum user object length attribute in the Root Quotas attributes page (see 7.1.3.12). For partition zero, the FORMAT OSD command (see 6.17) shall set the default maximum user object length attribute value to zero.

The capacity quota attribute (number 1 0001h) specifies the maximum value allowed in the used capacity attribute of the Partition Information attributes page (see 7.1.3.9). If the setting of an attribute value (see 5.2.6), an APPEND command (see 6.2), a CREATE command (see 6.5), a CREATE AND WRITE command (see 6.6), a CREATE COLLECTION command (see 6.8), or a WRITE command (see 6.40) attempts to exceed the capacity quota, a quota error (see 4.10.2) shall be generated. The CREATE PARTITION command shall set this attribute to the value in the partition capacity quota attribute in the Root Quotas attributes page.

The object count attribute (number 1 0002h) specifies the maximum value allowed in the number of collections and user objects attribute of the Partition Information attributes page. If a CREATE command (see 6.5), a CREATE AND WRITE command, or a CREATE COLLECTION command (see 6.8) attempts to exceed the object count quota, a quota error (see 4.10.2) shall be generated. The CREATE PARTITION command shall set this attribute to the value in the partition object count attribute in the Root Quotas attributes page. For partition zero, the FORMAT OSD command shall set the object count attribute value to zero.

The collections per user object (number 1 0081h) specifies the maximum number of collections in which a single user object may be a member. If a set attributes request specifying the Collections attributes page (see 7.1.3.21) attempts to exceed the collections per user object quota, a quota error (see 4.10.2) shall be generated. The CREATE PARTITION command shall set this attribute to the value in the partition collections per user object attribute in the Root Quotas attributes page. For partition zero, the FORMAT OSD command shall set the collections per user object attribute value to zero.

If a command attempts to set an attribute that table 175 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Partition Quotas attributes page is shown in table 176.

Table 176 — Partition Quotas attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) PAGE NUMBER (P+2h)							
3	(LSB)							
4	(MSB) PAGE LENGTH (1Ch)							
7	(LSB)							
8	(MSB) DEFAULT MAXIMUM USER OBJECT LENGTH							
15	(LSB)							
16	(MSB) CAPACITY QUOTA							
23	(LSB)							
24	(MSB) OBJECT COUNT							
31	(LSB)							
32	(MSB) COLLECTIONS PER USER OBJECT							
35	(LSB)							

The PAGE NUMBER field contains the attributes page number of the Partition Quotas attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Partition Quotas attributes page.

The DEFAULT MAXIMUM USER OBJECT LENGTH field contains the value of the default maximum user object length attribute.

The CAPACITY QUOTA field contains the value of the capacity quota attribute.

The OBJECT COUNT field contains the value of the object count attribute.

The COLLECTIONS PER USER OBJECT field contains the value of the collections per user object attribute.

7.1.3.14 User Object Quotas attributes page

The User Object Quotas attributes page (2h) shall contain the attributes listed in table 177. Except for attribute number 0h, all attributes in the User Object Quotas attributes page are quotas (see 4.10).

Table 177 — User Object Quotas attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	8	Maximum user object length	Yes	No
2h to FFFF FFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 User Object Quotas".

The maximum user object length attribute (number 1h) specifies the maximum value allowed in the user object logical length attribute of the User Object Information attributes page (see 7.1.3.11). If an APPEND command (see 6.2), a CREATE AND WRITE command (see 6.6), a WRITE command (see 6.40), or setting the user object logical length in the User Object Information attributes page (see 7.1.3.11) attempts to exceed the maximum user object length quota, a quota error (see 4.10.2) shall be generated. The CREATE command (see 6.5) and CREATE AND WRITE command shall set this attribute to the value in the default maximum user object length attribute in the Partition Quotas attributes page (see 7.1.3.13).

If a command attempts to set an attribute that table 177 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the User Object Quotas attributes page is shown in table 178.

Table 178 — User Object Quotas attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
3	PAGE NUMBER (2h)							(LSB)
4	(MSB)							
7	PAGE LENGTH (8h)							(LSB)
8	(MSB)							
15	MAXIMUM USER OBJECT LENGTH							(LSB)

The PAGE NUMBER field contains the attributes page number of the User Object Quotas attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the User Object Quotas attributes page.

The MAXIMUM USER OBJECT LENGTH field contains the value of the maximum user object length attribute.

7.1.3.15 Root Timestamps attributes page

The Root Timestamps attributes page (R+3h) shall contain the attributes listed in table 179. The updating of timestamp attributes in this page is controlled by the TIMESTAMPS CONTROL field (see 5.2.13) in the CDB.

Table 179 — Root Timestamps attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h		Reserved	No	
2h	6	Attributes accessed time	No	Yes
3h	6	Attributes modified time	No	Yes
4h to FFFF FFFDh		Reserved	No	
FFFF FFFEh	1	Timestamp bypass	Yes	No

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Root Timestamps".

The attributes accessed time attribute (number 2h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent command whose CDB get and set attributes parameters (see 5.2.6) transferred any attributes pages or values associated with the root object to the application client.

The attributes modified time attribute (number 3h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent command whose CDB get and set attributes parameters (see 5.2.6) set any attribute values associated with the root object.

The timestamp bypass attribute (number FFFF FFFEh) specifies the default timestamp update policy (see table 180) for the Root Timestamps page that is used under the control of the TIMESTAMPS CONTROL (see 5.2.13) field in the CDB.

Table 180 — Timestamp bypass attribute values

Value	Description
00h	Timestamps shall be updated as described in the subclause that defines them
01h to 7Eh	Reserved
7Fh	Timestamps shall not be updated ^a
80h to DFh	Reserved
E0 to FEh	Vendor specific
FFh	Timestamps shall be updated as specified by the TIMESTAMPS CONTROL field in the CDB (see 5.2.13)
^a A timestamp attribute that has never been updated shall have a length of six and a value of zero. Bypassing a timestamp update shall not affect any previously established timestamp attribute values.	

The FORMAT OSD command (see 6.17) shall set the timestamp bypass attribute to zero.

All commands received in the task set subsequent to the completion of a command that changes timestamp bypass attribute value shall be processed according to the new timestamp bypass attribute value. Each command in the task set concurrently with a command that changes the timestamp bypass attribute value may be processed with either the old or the new timestamp bypass attribute value in a vendor specific manner.

If a command attempts to set an attribute that table 179 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Root Timestamps attributes page is shown in table 181.

Table 181 — Root Timestamps attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	PAGE NUMBER (R+3h)						(LSB)
3								
4	(MSB)	PAGE LENGTH (Dh)						(LSB)
7								
8	(MSB)	ATTRIBUTES ACCESSED TIME						(LSB)
13								
14	(MSB)	ATTRIBUTES MODIFIED TIME						(LSB)
19								
20		TIMESTAMP BYPASS						

The PAGE NUMBER field contains the attributes page number of the Root Timestamps attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Root Timestamps attributes page.

The ATTRIBUTES ACCESSED TIME field contains the value of the attributes accessed time attribute.

The ATTRIBUTES MODIFIED TIME field contains the value of the attributes modified time attribute.

The TIMESTAMP BYPASS field contains the value of the timestamp bypass attribute.

7.1.3.16 Partition Timestamps attributes page

The Partition Timestamps attributes page (P+3h) shall contain the attributes listed in table 182. The updating of timestamp attributes in this page is controlled by the TIMESTAMPS CONTROL field (see 5.2.13) in the CDB.

Table 182 — Partition Timestamps attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	6	Created time	No	Yes
2h	6	Attributes accessed time	No	Yes
3h	6	Attributes modified time	No	Yes
4h	6	Data accessed time	No	Yes
5h	6	Data modified time	No	Yes
6h to FFFF FFFDh		Reserved	No	
FFFF FFFEh	1	Timestamp bypass	Yes	No

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Partition Timestamps".

For all partitions except partition zero, the created time attribute (number 1h) shall contain the value of the clock attribute in the Root Information attributes page (see 7.1.3.8) at the completion of the CREATE PARTITION command (see 6.8) that created the partition. For partition zero, the created time attribute shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent FORMAT OSD command (see 6.17).

The attributes accessed time attribute (number 2h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent command whose CDB get and set attributes parameters (see 5.2.6) transferred any attributes pages or values associated with the partition to the application client.

The attributes modified time attribute (number 3h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent command whose CDB get and set attributes parameters (see 5.2.6) set any attribute values associated with the partition.

For partition zero, the data accessed time attribute shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent LIST command that transferred a list of partitions in the root object to the application client. For all partitions except partition zero, the data accessed time attribute (number 4h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent:

- a) LIST command (see 6.20) that transferred a list of user objects in the partition to the application client; or
- b) LIST COLLECTION command (see 6.21) that transferred a list of collections in the partition to the application client.

For partition zero, the data modified time attribute shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent CREATE PARTITION command (see 6.9) or REMOVE PARTITION command (see 6.34) that created or removed a partition. For all partitions except partition zero, the

data modified time attribute (number 5h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent CREATE command (see 6.5), CREATE AND WRITE command (see 6.6), CREATE COLLECTION command (see 6.8), REMOVE command (see 6.31), or REMOVE COLLECTION command (see 6.32) that created or removed a collection or user object in the partition.

The timestamp bypass attribute (number FFFF FFFEh) specifies the default timestamp update policy (see table 180 in 7.1.3.15) that is used for the following timestamp attributes pages used under the control of the TIMESTAMPS CONTROL (see 5.2.13) field in the CDB:

- a) Partition Timestamps page;
- b) Collection Timestamps attributes page (see 7.1.3.17); and
- c) User Object Timestamps attributes page (see 7.1.3.18).

All commands received in the task set subsequent to the completion of a command that changes timestamp bypass attribute value shall be processed according to the new timestamp bypass attribute value. Each command in the task set concurrently with a command that changes the timestamp bypass attribute value may be processed with either the old or the new timestamp bypass attribute value in a vendor specific manner.

The CREATE PARTITION command (see 6.9) shall set this attribute to the value in the timestamp bypass attribute in the Root Timestamps attributes page (see 7.1.3.15).

If a command attempts to set an attribute that table 182 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Partition Timestamps attributes page is shown in table 183.

Table 183 — Partition Timestamps attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
3	PAGE NUMBER (P+3h)							(LSB)
4	(MSB)							
7	PAGE LENGTH (1Fh)							(LSB)
8	(MSB)							
13	CREATED TIME							(LSB)
14	(MSB)							
19	ATTRIBUTES ACCESSED TIME							(LSB)
20	(MSB)							
25	ATTRIBUTES MODIFIED TIME							(LSB)
26	(MSB)							
31	DATA ACCESSED TIME							(LSB)
32	(MSB)							
37	DATA MODIFIED TIME							(LSB)
38	TIMESTAMP BYPASS							

The PAGE NUMBER field contains the attributes page number of the Partition Timestamps attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Partition Timestamps attributes page.

The CREATED TIME field contains the value of the created time attribute.

The ATTRIBUTES ACCESSED TIME field contains the value of the attributes accessed time attribute.

The ATTRIBUTES MODIFIED TIME field contains the value of the attributes modified time attribute.

The DATA ACCESSED TIME field contains the value of the data accessed time attribute.

The DATA MODIFIED TIME field contains the value of the data modified time attribute.

The TIMESTAMP BYPASS field contains the value of the timestamp bypass attribute.

7.1.3.17 Collection Timestamps attributes page

The Collection Timestamps attributes page (C+3h) shall contain the attributes listed in table 184. The updating of timestamp attributes in this page is controlled by the TIMESTAMPS CONTROL field (see 5.2.13) in the CDB.

Table 184 — Collection Timestamps attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	6	Created time	No	Yes
2h	6	Attributes accessed time	No	Yes
3h	6	Attributes modified time	No	Yes
4h	6	Data accessed time	No	Yes
5h	6	Data modified time	No	Yes
6h to FFFF FFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Collection Timestamps".

The created time attribute (number 1h) shall contain the value of the clock attribute in the Root Information attributes page (see 7.1.3.8) at the completion of the CREATE COLLECTION command (see 6.8) that created the associated collection.

The attributes accessed time attribute (number 2h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent command whose CDB get and set attributes parameters (see 5.2.6) transferred any attributes pages or values associated with the collection to the application client.

The attributes modified time attribute (number 3h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent command whose CDB get and set attributes parameters (see 5.2.6) set any attribute values associated with the collection.

The data accessed time attribute (number 4h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent LIST COLLECTION command (see 6.21) that transferred a list of user objects in the collection to the application client.

The data modified time attribute (number 5h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent:

- a) Set attributes command function to a user object Collections attributes page (see 7.1.3.21) that added or removed a member from the collection; or
- b) Multi-object command (see 4.6.6.6) that removed a user object from the collection.

If a command attempts to set an attribute that table 184 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Collection Timestamps attributes page is shown in table 185.

Table 185 — Collection Timestamps attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	PAGE NUMBER (C+3h)						(LSB)
3								
4	(MSB)	PAGE LENGTH (1Eh)						(LSB)
7								
8	(MSB)	CREATED TIME						(LSB)
13								
14	(MSB)	ATTRIBUTES ACCESSED TIME						(LSB)
19								
20	(MSB)	ATTRIBUTES MODIFIED TIME						(LSB)
25								
26	(MSB)	DATA ACCESSED TIME						(LSB)
31								
32	(MSB)	DATA MODIFIED TIME						(LSB)
37								

The PAGE NUMBER field contains the attributes page number of the Collection Timestamps attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Collection Timestamps attributes page.

The CREATED TIME field contains the value of the created time attribute.

The ATTRIBUTES ACCESSED TIME field contains the value of the attributes accessed time attribute.

The ATTRIBUTES MODIFIED TIME field contains the value of the attributes modified time attribute.

The DATA ACCESSED TIME field contains the value of the data accessed time attribute.

The DATA MODIFIED TIME field contains the value of the data modified time attribute.

7.1.3.18 User Object Timestamps attributes page

The User Object Timestamps attributes page (3h) shall contain the attributes listed in table 186. The updating of timestamp attributes in this page is controlled by the TIMESTAMPS CONTROL field (see 5.2.13) in the CDB.

Table 186 — User Object Timestamps attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	6	Created time	No	Yes
2h	6	Attributes accessed time	No	Yes
3h	6	Attributes modified time	No	Yes
4h	6	Data accessed time	No	Yes
5h	6	Data modified time	No	Yes
6h to FFFF FFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 User Object Timestamps".

The created time attribute (number 1h) shall contain the value of the clock attribute in the Root Information attributes page (see 7.1.3.8) at the completion of the CREATE command (see 6.5) or CREATE AND WRITE command (see 6.6) that created the associated user object.

The attributes accessed time attribute (number 2h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent:

- a) Command whose CDB get and set attributes parameters (see 5.2.6) transferred any attributes pages or values associated with the user object to the application client, or
- b) QUERY command (see 6.26) that evaluated any attributes values associated with the user object to the application client during attempts to match the values specified in the query list.

The attributes modified time attribute (number 3h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent command whose CDB get and set attributes parameters (see 5.2.6) set any attribute values associated with the user object.

The data accessed time attribute (number 4h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent READ command (see 6.27) that transferred data from the user object to the application client.

The data modified time attribute (number 5h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent command that changed the value of the user object logical length attribute in the User Object Information attributes page (see 7.1.3.11) or that stored data in the user object (i.e., a WRITE command (see 6.40), an APPEND command (see 6.2), or a CREATE AND WRITE command (see 6.6)).

If a command attempts to set an attribute that table 186 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the User Object Timestamps attributes page is shown in table 187.

Table 187 — User Object Timestamps attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	PAGE NUMBER (3h)						(LSB)
3								
4	(MSB)	PAGE LENGTH (1Eh)						(LSB)
7								
8	(MSB)	CREATED TIME						(LSB)
13								
14	(MSB)	ATTRIBUTES ACCESSED TIME						(LSB)
19								
20	(MSB)	ATTRIBUTES MODIFIED TIME						(LSB)
25								
26	(MSB)	DATA ACCESSED TIME						(LSB)
31								
32	(MSB)	DATA MODIFIED TIME						(LSB)
37								

The PAGE NUMBER field contains the attributes page number of the User Object Timestamps attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the User Object Timestamps attributes page.

The CREATED TIME field contains the value of the created time attribute.

The ATTRIBUTES ACCESSED TIME field contains the value of the attributes accessed time attribute.

The ATTRIBUTES MODIFIED TIME field contains the value of the attributes modified time attribute.

The DATA ACCESSED TIME field contains the value of the data accessed time attribute.

The DATA MODIFIED TIME field contains the value of the data modified time attribute.

7.1.3.19 Attributes Access attributes page

The Attributes Access attributes page (P+4h) shall contain the attributes listed in table 188.

Table 188 — Attributes Access attributes page contents

Attribute Number	Length (bytes) ^a	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h to FFFF FFFEh	0 or n	Allowed attributes access	Yes	No

^a A length of 0 in this column denotes an attribute that may be undefined (see 3.1.51).

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Attributes Access".

Each allowed attributes access attribute (1h to FFFF FFFEh) is:

- a) An undefined (i.e., zero length) attribute (see 3.1.51); or
- b) A defined attribute (see 3.1.14) that contains a list of attributes access descriptors (see table 189) each of which indicates an attribute or attributes to which access is allowed.

Table 189 — Allowed attributes access attribute format

Bit Byte	7	6	5	4	3	2	1	0
	Attributes access descriptors							
0	Attributes access descriptor [first] (see table 190)							
7								
	⋮							
n-7	Attributes access descriptor [last] (see table 190)							
n								

Each attributes access descriptor (see table 190) indicates an attribute or set of attributes to which this descriptor allows access.

Table 190 — Attributes access descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) ATTRIBUTES PAGE (LSB)							
3								
4	(MSB) ATTRIBUTE NUMBER (LSB)							
7								

The ATTRIBUTES PAGE field identifies the page number of one attribute or set of attributes to which access is allowed.

The ATTRIBUTE NUMBER field identifies:

- a) The attribute number within the attributes page specified by the ATTRIBUTES PAGE field of the one attribute to which access is allowed; or
- b) The value FFFF FFFFh indicates that access is allowed to all the attributes in the attributes page specified by the ATTRIBUTES PAGE field.

If the ALLOWED ATTRIBUTES ACCESS field in a capability (see 4.11.2.2) specifies the number of a defined attribute in an Attributes Access attributes page and the command attempts to retrieve or set an attribute that is not identified in at least one attributes access descriptor, then the command shall be terminated as described in 7.1.2.

An Attributes Access attributes page allowed access attribute serves only to disallow access to attributes to which access would otherwise be possible and this is accomplished by omitting them from the attributes identified in the

attributes access descriptors. The Attributes Access attributes page does not grant access to attributes that a command would not otherwise be able to access.

If a command attempts to set an attribute that table 188 states is not application client settable, then the command shall be terminated as described in 7.1.2.

7.1.3.20 Command Tracking attributes page

The Command Tracking attributes page (C+4h) shall contain the attributes listed in table 191.

Table 191 — Command Tracking attributes page contents

Attribute Number	Length (bytes) ^a	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	1	Percent complete	No	Yes
2h	2	Active command status	No	Yes
3h	2	Ended command status	No	Yes
4h	0 or n	Sense data	No	Yes
5h to Fh		Reserved	No	
10h	8	Number of members	No	Yes
11h	0 or 8	Objects processed	No	Yes
12h	0 or 8	Newer objects skipped	No	Yes
13h	0 or 8	Missing objects skipped	No	Yes
14h to EFFF FFFFh		Reserved	No	
F000 0000h to FFFF FFFEh	0 or n	Vendor specific ^b	No	Yes

^a A length of 0 in this column denotes an attribute that may be undefined (see 3.1.51).

^b The combination of a TRACKING collection's (see 4.6.6.3) membership and the Command Tracking attributes page attributes shall be sufficient to restart an interrupted command (e.g., an interrupted CREATE SNAPSHOT command (see 6.10)) or a command that was terminated with CHECK CONDITION status. Information in the vendor specific attributes may be needed to fulfill this requirement.

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Command Tracking".

The percent complete attribute (number 1h) shall indicate percentage of the processing that has been completed for the command, if any, for which the device server is using the collection to track processing activities.

The active command status attribute (number 2h) shall indicate the nature of the command (see table 192), if any, that is using the TRACKING collection to track processing activities.

Table 192 — Active command status attribute values

Active command status attribute value	Description
0000h	No command is using the collection to track processing activities.
0001h to 87FFh	Reserved
8800h to 8FFFFh	A command whose SERVICE ACTION field contains the attribute value is using the collection to track processing activities.
9000h to FFFFh	Reserved

The ended command status attribute (number 3h) shall indicate the status (see table 193) of the most recent command that used the TRACKING collection that has ended.

Table 193 — Ended command status attribute values

Ended command status attribute value	Description
0000h	Command processing has completed with GOOD status.
0001h to 00FFh	Command processing has completed with the status code (see SAM-4) contained in the least significant byte of the attribute value.
0100h to 7FFFh	Reserved
8000h	The command that was using the collection to track processing activities was interrupted for an unknown reason.
8001h	The command that was using the collection to track processing activities was interrupted an ABORT TASK task management function (see SAM-4) or another condition whose processing emulates an ABORT TASK task management function.
8002h	The command that was using the collection to track processing activities was interrupted by a power on event (See SAM-4).
8003h	The command that was using the collection to track processing activities was interrupted by a bus reset event (See SAM-4).
8004h	The command that was using the collection to track processing activities was interrupted by a logical unit reset event (See SAM-4).
8005h	The command that was using the collection to track processing activities was interrupted by a I_T nexus loss event (See SAM-4).
8006h	The command that was using the collection to track processing activities was interrupted by a power loss expected event (See SAM-4).
8007h to FFFEh	Reserved
FFFFh	The ending status of a command that has used the collection to track processing activities, if any, is not available.

If the ended command status attribute is set to 0002h (i.e., command processing has completed with CHECK CONDITION status), then the sense data attribute (number 4h) shall contain the sense data that was, or should have been, returned to the application client. If the ended command status attribute is not set to 0002h, the sense data attribute should be undefined (see 3.1.51).

The number of members attribute (number 10h) shall indicate the number of objects that are members of the collection.

If it is defined (see 3.1.14), the objects processed attribute shall indicate the number of objects that have been removed from the collection following successful processing as specified by the command (e.g., as described for multi-object commands in 4.6.6.6).

If it is defined (see 3.1.14), the newer objects skipped attribute shall indicate the number of objects that have been removed from the collection because the creation time attribute in the User Object Timestamps attributes page (see 7.1.3.18) is later than (i.e., greater than) the creation time attribute in the Collection Timestamps attributes page (see 7.1.3.17) (e.g., as described for multi-object commands in 4.6.6.6).

If it is defined (see 3.1.14), the missing objects skipped attribute shall indicate the number of objects that have been removed from the collection because the object was not present in the partition at the time processing was attempted (e.g., as described for multi-object commands in 4.6.6.6).

If a command attempts to set an attribute that table 191 states is not application client settable, then the command shall be terminated as described in 7.1.2.

7.1.3.21 Collections attributes page

The Collections attributes page (4h) shall contain the attributes listed in table 194.

Table 194 — Collections attributes page contents

Attribute Number	Length (bytes) ^a	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	0 or 40	Page identification	No	Yes
1h to FFFF FF00h	0 or 8	Collection pointer	Yes/No	No
FFFF FF01h to FFFF FFEh		Reserved	No	
^a A length of 0 in this column denotes an attribute that may be undefined (see 3.1.51).				

If collections are supported, the page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Collections". If collections are not supported, the length of the page identification attribute shall be zero.

Each collection pointer attribute (1h to FFFF FF00h) may be:

- a) A zero length attribute (i.e., contain no value); or
- b) The Collection_Object_ID of a collection (see 4.6.6) to which the user object belongs.

For a LINKED collection (see 4.6.6.2), a user object is made a member of a collection by setting one of its collection pointer attribute values to the Collection_Object_ID of that collection.

For a LINKED collection (see 4.6.6.2), a user object is removed from the membership of a collection by:

- a) Changing the collection pointer attribute identifying that collection to have a length of zero; or
- b) Setting the collection pointer attribute identifying that collection to the Collection_Object_ID of a different collection.

The command shall be terminated as described in 7.1.2 if it attempts to set:

- a) The same Collection_Object_ID in more than one collection pointer attribute;
- b) A collection pointer attribute to a value that is not a Collection_Object_ID;
- c) A collection pointer attribute to the Collection_Object_ID of a collection in which the collection type attribute in the Collection Information attributes page (see 7.1.3.10) is set to a value other than LINKED (see table 171 in 7.1.3.10); or
- d) A collection pointer attribute to any length other than zero or eight.

If setting a collection pointer attribute causes the number of collection pointer attributes with non-zero attribute lengths to exceed the value in the collections per user object attribute in the Partition Quotas attributes page (see 7.1.3.13), then a quota error shall be generated (see 4.10.2). The quota testing principles described in 4.10.3 apply to the testing of the object count quota.

If a command attempts to set an attribute that table 194 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Collections attributes page is shown in table 195.

Table 195 — Collections attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	PAGE NUMBER (4h)						(LSB)
3								
4	(MSB)	PAGE LENGTH (n-7)						(LSB)
7								
8	(MSB)	First collection pointer attribute value						(MSB)
15								
16	(MSB)	Second collection pointer attribute value						(MSB)
23								
		⋮						
n-7	(MSB)	Last collection pointer attribute value						(MSB)
n								

The PAGE NUMBER field contains the attributes page number of the Collection attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Collection attributes page.

The first collection pointer attribute value shall contain the attribute value for the lowest numbered collection pointer attribute with a length of eight.

The second collection pointer attribute value shall contain the attribute value for the second lowest numbered collection pointer attribute with a length of eight.

Additional collection pointer attribute values shall be added to the page format for each collection pointer attribute with a length of eight.

The last collection pointer attribute value shall contain the attribute value for the highest numbered collection pointer attribute with a length of eight.

7.1.3.22 Root Policy/Security attributes page

The Root Policy/Security attributes page (R+5h) shall contain the attributes listed in table 196.

Table 196 — Root Policy/Security attributes page contents

Attribute Number	Length (bytes) ^a	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	1	Default security method	Yes	Yes
2h	6	Oldest valid nonce limit	No	Yes
3h	6	Newest valid nonce limit	No	Yes
4h to 5h		Reserved	No	
6h	1	Partition default security method	Yes	Yes
7h	2	Supported security methods	No	Yes
8h		Reserved	No	
9h	6	Adjustable clock	Yes	Yes
Ah	2	Boot epoch	Yes	Yes
Bh to 7FFCh		Reserved	No	
7FFDh	0 or 7	Master key identifier	No	Yes
7FFEh	0 or 7	Root key identifier	No	Yes
7FFFh to 7FFF FFFFh		Reserved	No	
8000 0000h to 8000 00Fh	1	Supported integrity check value algorithm	No	Yes
8000 0010h to 8000 001Fh	1	Supported DH group	No	Yes
8000 0020h to FFFF FFFEh		Reserved	No	

^a A length of 0 in this column denotes an attribute that may be undefined (see 3.1.51).

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Root Policy/Security".

The default security method attribute (number 1h) specifies the security method (see 4.12.4) used for the processing of the SET KEY command (see 6.37) and SET MASTER KEY command (see 6.38) in the absence of conditions that specify a different security method (see 4.12.3). The value of the default security method attribute shall not be changed by a FORMAT OSD command (see 6.17). The value placed in the default security method attribute when the OBSD (see 3.1.27) is manufactured is vendor specific. If the value of the default security method attribute is changed, the application client should invalidate the working keys for partition zero using the SET KEY command.

The oldest valid nonce limit attribute (number 2h) specifies the largest value allowed in the oldest valid nonce attribute in any Partition Policy/Security attributes page (see 7.1.3.23) (i.e., the maximum number of milliseconds prior to the value in the clock attribute in the Root Information attributes page (see 7.1.3.8) to which the device server constrains the contents of the **TIMESTAMP** field in a request nonce (see 4.9.6)).

The newest valid nonce limit attribute (number 3h) specifies the largest value allowed in the newest valid nonce attribute in any Partition Policy/Security attributes page (i.e., the maximum number of milliseconds subsequent to the value in the clock attribute in the Root Information attributes page to which the device server constrains the contents of the **TIMESTAMP** field in a request nonce).

The partition default security method attribute (number 6h) specifies the value to be placed in the default security method attribute of each partition, when it is created. The value of the partition default security method attribute shall not be changed by a **FORMAT OSD** command (see 6.17). The value placed in the partition default security method attribute when the **OBSD** is manufactured is vendor specific.

The supported security methods attribute (number 7h) indicates which security methods (see 4.12.4) are supported by the **OSD** logical unit (see table 197).

Table 197 — Supported security methods attribute format

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved				ALLDATA	CMDRSP	CAPKEY	NOSEC
1	Reserved							

The **NOSEC** (**NOSEC** security method supported) bit is set to zero if the **NOSEC** security method is not supported. The **NOSEC** bit is set to one if the **NOSEC** security method is supported.

The **CAPKEY** (**CAPKEY** security method supported) bit is set to zero if the **CAPKEY** security method is not supported. The **CAPKEY** bit is set to one if the **CAPKEY** security method is supported.

The **CMDRSP** (**CMDRSP** security method supported) bit is set to zero if the **CMDRSP** security method is not supported. The **CMDRSP** bit is set to one if the **CMDRSP** security method is supported.

The **ALLDATA** (**ALLDATA** security method supported) bit is set to zero if the **ALLDATA** security method is not supported. The **ALLDATA** bit is set to one if the **ALLDATA** security method is supported.

The adjustable clock attribute (number 9h) shall contain the current time in use by the **OSD** device server represented as the count of the number of milliseconds elapsed since midnight, 1 January 1970 UT (see 3.1.52). The value shall be set to the UT when the **OBSD** (see 3.1.27) is manufactured and may be modified by the application client after that. The mechanism used to maintain the adjustable clock attribute value is outside the scope of the standard. The adjustable clock attribute value should not gain or lose more than one second in any 24-hour interval.

When the **OSD** device is manufactured, the boot epoch attribute (number Ah) should be set to a non-zero value. The processing of a power on SCSI device condition, hard reset SCSI device condition, or logical unit reset SCSI device condition established in response to an event (see SAM-4) shall cause the boot epoch attribute to be updated as follows:

- a) If the boot epoch attribute is not set to FFFFh, then one shall be added to the boot epoch value; or
- b) If the boot epoch attribute is set to FFFFh, then the boot epoch value shall be set to one.

The boot epoch attribute is compared to the **BOOT EPOCH** field in capabilities as described in 4.11.2.2.

NOTE 8 The application client may change the boot epoch attribute to any value (see table 196).

The master key identifier attribute (number 7FFDh) contains the key identifier value from the most recent successful SET MASTER KEY command (see 6.38). If a SET MASTER KEY command has never been processed, the master key identifier attribute length shall be seven and the master key identifier attribute value shall be the ASCII characters "1st key".

The root key identifier attribute (number 7FFEh) contains the key identifier value from the most recent successful SET KEY command (see 6.37) with the KEY TO SET field set to 01b (i.e., update root key). If the root key is invalid (i.e., never set or invalidated by a SET MASTER KEY command), the root key identifier attribute length shall be zero. Regardless of the root key identifier attribute length, the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) for partition zero (see 3.1.33) shall reflect an attribute length of seven (i.e., it shall not be possible for a SET KEY command to cause the partition zero used capacity attribute value to exceed the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13) for partition zero and generate a quote error).

The supported integrity check value algorithm attributes (numbers 8000 0000h to 8000 000Fh) contain coded values (see table 198) identifying the supported algorithms that the OSD device server supports for computing integrity check values.

The supported integrity check value algorithm with the lowest valued attribute number (i.e., 8000 0000h) identifies the most preferred integrity check value algorithm and the highest valued attribute number (i.e., 8000 000Fh) identifies the least preferred algorithm. If a supported integrity check value algorithm attribute contains zero, then all supported integrity check value algorithm attributes with higher valued attribute numbers also shall contain zero.

The low order four bits of the attribute number are the value that appears in the capability INTEGRITY CHECK VALUE ALGORITHM field (see 4.11.2.2) in each capability (e.g., attribute number 8000 0007h identifies the integrity check value algorithm used if the INTEGRITY CHECK VALUE ALGORITHM field contains seven).

Table 198 — Supported integrity check value algorithm codes

Value	Algorithm	Reference
00h	No algorithm supported	FIPS 180-1 (1995) and FIPS 198 (2002)
01h	HMAC-SHA1	
02h to DFh	Reserved	
E0h to FFh	Vendor specific	

The supported DH group attributes (numbers 8000 0010h to 8000 001Fh) contain coded values identifying the supported values in the DH_GROUP field of a SET MASTER KEY command (see 6.38). The values of the supported DH group attributes are the values associated with the Group Description class (i.e., class code value 4) in the Internet Key Exchange Attributes registry maintained by IANA (see <http://www.iana.org/assignments/ipsec-registry>). The DH group indicated by each value is as specified by IANA in that registry.

Every DH group identified by a supported DH group attribute shall be a MODP DH group. The code values 1h (i.e., the 768-bit MODP DH group defined by RFC 2409) and 2h (i.e., the 1024-bit MODP DH group defined by RFC 2409) shall not appear in any supported DH group attribute.

NOTE 9 The constraint to MODP DH groups eliminates usage of all elliptic curve DH groups (e.g., the DH groups having code values 3, 4, and 6 through 13, inclusive).

One of the supported DH group attributes shall contain Dh (i.e., 14) indicating the 2048-bit MODP DH group defined by RFC 3526.

The supported DH group with the lowest valued attribute number (i.e., 8000 0000h) identifies the most preferred DH group and the highest valued attribute number (i.e., 8000 000Fh) identifies the least preferred DH group. If a supported DH group attribute contains zero, then all supported DH group attributes with higher valued attribute numbers also shall contain zero

If a command attempts to set an attribute that table 196 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Root Policy/Security attributes page is shown in table 199.

Table 199 — Root Policy/Security attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	PAGE NUMBER (R+5h)						(LSB)
3								
4	(MSB)	PAGE LENGTH (42h)						(LSB)
7								
8		DEFAULT SECURITY METHOD						
9		PARTITION DEFAULT SECURITY METHOD						
10		SUPPORTED SECURITY METHODS						
11								
12	(MSB)	OLDEST VALID NONCE LIMIT						(LSB)
17								
18	(MSB)	NEWEST VALID NONCE LIMIT						(LSB)
23								
24		Reserved				MKI_VALID		RKI_VALID
25	(MSB)	MASTER KEY IDENTIFIER						(LSB)
31								
32	(MSB)	ROOT KEY IDENTIFIER						(LSB)
38								
39		Most preferred SUPPORTED INTEGRITY CHECK VALUE ALGORITHM (attribute number 8000 0000h)						
		⋮						
54		Least preferred SUPPORTED INTEGRITY CHECK VALUE ALGORITHM (attribute number 8000 000Fh)						
55		Most preferred SUPPORTED DH_GROUP (attribute number 8000 0010h)						
		⋮						
70		Least preferred SUPPORTED DH_GROUP (attribute number 8000 001Fh)						
71		Reserved						
72	(MSB)	BOOT EPOCH						(LSB)
73								

The PAGE NUMBER field contains the attributes page number of the Root Policy/Security attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Root Policy/Security attributes page.

The DEFAULT SECURITY METHOD field contains the value of the default security method attribute.

The PARTITION DEFAULT SECURITY METHOD field contains the value of the partition default security method attribute.

The SUPPORTED SECURITY METHODS field contains the value of the supported security methods attribute.

The OLDEST VALID NONCE LIMIT field contains the value of the oldest valid nonce limit attribute.

The NEWEST VALID NONCE LIMIT field contains the value of the newest valid nonce limit attribute.

The MKI_VALID (master key identifier valid) bit shall be set to zero if the master key identifier attribute length is zero. Otherwise, the MKI_VALID bit shall be set to one.

The RKI_VALID (root key identifier valid) bit shall be set to zero if the root key identifier attribute length is zero. Otherwise, the RKI_VALID bit shall be set to one.

If the MKI_VALID bit is set to one, the MASTER KEY IDENTIFIER field contains the value of the master key identifier attribute. Otherwise, the contents of the MASTER KEY IDENTIFIER field are undefined.

If the RKI_VALID bit is set to one, the ROOT KEY IDENTIFIER field contains the value of the root key identifier attribute. Otherwise, the contents of the ROOT KEY IDENTIFIER field are undefined.

The sixteen SUPPORTED INTEGRITY CHECK VALUE ALGORITHM fields contain the supported integrity check value attribute values in ascending attribute number order. The SUPPORTED INTEGRITY CHECK VALUE ALGORITHM field with the smallest byte offset in the page identifies the most preferred integrity check value algorithm. The SUPPORTED INTEGRITY CHECK VALUE ALGORITHM field with the largest byte offset in the page identifies the least preferred algorithm.

The sixteen SUPPORTED DH GROUP fields contain the supported DH group attribute values in ascending attribute number order. The SUPPORTED DH GROUP field with the smallest byte offset in the page identifies the most preferred DH group to be used by the SET MASTER KEY command (see 6.38). The SUPPORTED DH GROUP field with the largest byte offset in the page identifies the least preferred DH group.

The BOOT EPOCH field contains the value of the boot epoch attribute.

7.1.3.23 Partition Policy/Security attributes page

The Partition Policy/Security attributes page (P+5h) shall contain the attributes listed in table 200.

Table 200 — Partition Policy/Security attributes page contents

Attribute Number	Length (bytes) ^a	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	1	Default security method	Yes	Yes
2h	6	Oldest valid nonce	Yes	Yes
3h	6	Newest valid nonce	Yes	Yes
4h	2	Request nonce list depth	No	Yes
5h	2	Frozen working key bit mask	No	Yes
6h to 7FFEh		Reserved	No	
7FFFh	0 or 7	Partition key identifier	No	Yes
8000h to 800Fh	0 or 7	Working key identifier	No	Yes
8010h to 4000 0000h		Reserved	No	
4000 0001h	4	Policy access tag	Yes	Yes
4000 0002h	4	User object policy access tag	Yes	Yes
4000 0003h to FFFF FFFEh		Reserved	No	
^a A length of 0 in this column denotes an attribute that may be undefined (see 3.1.51).				

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Partition Policy/Security".

The default security method attribute (number 1h) specifies the security method (see 4.12.4) used for the processing of all commands except the SET KEY command and SET MASTER KEY command in the absence of conditions that specify a different security method (see 4.12.3). The value of the default security method attribute for partition zero shall not be changed by a FORMAT OSD command (see 6.17). The value placed in the default security method attribute for partition zero when the OBSD (see 3.1.27) is manufactured is vendor specific. If the value of the default security method attribute is changed, the working keys for affected partition should be invalidated using the SET KEY command (see 6.37).

A CREATE PARTITION command (see 6.9) shall copy the partition default security method attribute from the Root Policy/Security attributes page (see 7.1.3.22) to the default security method attribute in new Partition Policy/Security attributes page.

The oldest valid nonce attribute (number 2h) indicates the number of milliseconds prior to the value in the clock attribute in the Root Information attributes page (see 7.1.3.8) to which the device server constrains the contents of the TIMESTAMP field in a request nonce (see 4.12.7) received in a command addressed to the partition, a collection in the partition, or a user object in the partition. The processing of request nonces affected by this constraint is described in 4.12.7.2.

If a command contains a request to set the oldest valid nonce attribute to a value that is larger than the value in the oldest valid nonce limit attribute in the Root Policy/Security attributes page (see 7.1.3.22), then the command shall be terminated as described in 7.1.2.

The newest valid nonce attribute (number 3h) indicates the number of milliseconds later than the value in the clock attribute in the Root Information attributes page to which the device server constrains the contents of the TIMESTAMP field in a request nonce (see 4.12.7) received in a command addressed to the partition, a collection in the partition, or a user object in the partition. The processing of request nonces affected by this constraint is described in 4.12.7.2.

If a command contains a request to set the newest valid nonce attribute to a value that is larger than the value in the newest valid nonce limit attribute in the Root Policy/Security attributes page, then the command shall be terminated as described in 7.1.2.

The request nonce list depth attribute (number 4h) shall contain the minimum number of request nonce list entries 4.12.7.3.1 available to one application client.

The frozen working key bit mask attribute (number 5h) indicates which working key versions (see table 201) have been frozen as part of request nonce list processing (see 4.12.7.3.3).

Table 201 — Frozen working key bit mask attribute format

Bit Byte	7	6	5	4	3	2	1	0
0	WK07_FZN	WK06_FZN	WK05_FZN	WK04_FZN	WK03_FZN	WK02_FZN	WK01_FZN	WK00_FZN
1	WK0F_FZN	WK0E_FZN	WK0D_FZN	WK0C_FZN	WK0B_FZN	WK0A_FZN	WK09_FZN	WK08_FZN

A WK00_FZN (working key 0h frozen) bit set to zero indicates that device server is not rejecting commands that contain capabilities with the working key with a key version of zero as part of request nonce list processing. A WK00_FZN bit set to one indicates that device server is rejecting commands that contain capabilities with the working key with a key version of zero as part of request nonce list processing (see 4.12.7.3.3). Once the WK00_FZN bit is set to one, it shall not be set to zero until a new working key with key version zero is established using the SET KEY command (see 6.37).

The WK01_FZN bit, WK01_FZN bit, WK02_FZN bit, WK03_FZN bit, WK04_FZN bit, WK05_FZN bit, WK06_FZN bit, WK07_FZN bit, WK08_FZN bit, WK09_FZN bit, WK0A_FZN bit, WK0B_FZN bit, WK0C_FZN bit, WK0D_FZN bit, WK0E_FZN bit, and WK0F_FZN have the same bit value definitions as the WK00_FZN bit, except that the definitions apply to the working keys with key versions one to fifteen, respectively.

The partition key identifier attribute (number 7FFFh) contains the key identifier value from the most recent successful SET KEY command (see 6.37) with the KEY TO SET field set to 10b (i.e., update partition key). If the partition key is invalid (i.e., never set, invalidated by a SET MASTER KEY command (see 6.38), or invalidated by a SET KEY command), the partition key identifier attribute length shall be zero. Regardless of the partition key identifier attribute length, the used capacity attribute in the Partition Information attributes page (see 7.1.3.9) shall reflect an attribute length of seven (i.e., it shall not be possible for a SET KEY command to cause the partition's used capacity attribute value to exceed the capacity quota attribute in the Partition Quotas attributes page (see 7.1.3.13) and generate a quote error).

The working key identifier attributes (numbers 8000h to 800Fh) contain the key identifier value from the most recent successful SET KEY command with:

- a) The KEY TO SET field set to 11b (i.e., update working key); and
- b) The KEY VERSION field set to the attribute number minus 8000h (e.g., a version key of three sets attribute 8003h and a version key of eight sets attribute 8008h).

If a working key is invalid (i.e., never set, invalidated by a SET MASTER KEY command, or invalidated by a SET KEY command), the working key identifier attribute length for the associated working key shall be zero. Regardless of the lengths of any of the working key identifier attributes, the used capacity attribute in the Partition Information attributes page shall reflect an attribute length of seven for all sixteen working key identifier attributes (i.e., it shall not be possible for a SET KEY command to cause the partition's used capacity attribute value to exceed the capacity quota attribute in the Partition Quotas attributes page and generate a quote error).

The policy access tag attribute (number 4000 0001h) specifies the expected non-zero contents of the POLICY ACCESS TAG field in any capability (see 4.11.2) that allows access to this partition. The format, use, and attribute setting restrictions for the policy access tag attribute are described in 4.11.3.2. A CREATE PARTITION command (see 6.9) shall set the policy access tag attribute to 7FFF FFFFh.

The user object policy access tag attribute (number 4000 0002h) specifies the value to be placed in the policy access tag attribute of each collection or user object, when it is created. A CREATE PARTITION command shall set the user object policy access tag attribute to 7FFF FFFFh.

If a command attempts to set an attribute that table 200 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Partition Policy/Security attributes page is shown in table 202.

Table 202 — Partition Policy/Security attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
3	PAGE NUMBER (P+5h)							(LSB)
4	(MSB)							
7	PAGE LENGTH (8Fh)							(LSB)
8								
10	Reserved							
11	DEFAULT SECURITY METHOD							
12	(MSB)							
17	OLDEST VALID NONCE							(LSB)
18	(MSB)							
23	NEWEST VALID NONCE							(LSB)
24	(MSB)							
25	REQUEST NONCE LIST DEPTH							(LSB)
26								
27	FROZEN WORKING KEY BIT MASK							
28	(MSB)							
31	POLICY ACCESS TAG							(LSB)
32	(MSB)							
35	USER OBJECT POLICY ACCESS TAG							(LSB)
36	Reserved							PKI_VALID
37	WKI07_VLD	WKI06_VLD	WKI05_VLD	WKI04_VLD	WKI03_VLD	WKI02_VLD	WKI01_VLD	WKI00_VLD
38	WKI0F_VLD	WKI0E_VLD	WKI0D_VLD	WKI0C_VLD	WKI0B_VLD	WKI0A_VLD	WKI09_VLD	WKI08_VLD
39	(MSB)							
45	PARTITION KEY IDENTIFIER							(LSB)
46	(MSB)							
64	WORKING KEY IDENTIFIER (for attribute number 8000h)							(LSB)
	⋮							
144	(MSB)							
150	WORKING KEY IDENTIFIER (for attribute number 800Fh)							(LSB)

The PAGE NUMBER field contains the attributes page number of the Partition Policy/Security attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Partition Policy/Security attributes page.

The DEFAULT SECURITY METHOD field contains the value of the default security method attribute.

The OLDEST VALID NONCE field contains the value of the oldest valid nonce attribute.

The NEWEST VALID NONCE field contains the value of the newest valid nonce attribute.

The REQUEST NONCE LIST DEPTH field contains the value of the request nonce list depth attribute.

The FROZEN WORKING KEY BIT MASK field contains the value of the frozen working key bit mask attribute.

The POLICY ACCESS TAG field contains the value of the policy access tag attribute.

The USER OBJECT POLICY ACCESS TAG field contains the value of the user object policy access tag attribute.

The PKI_VALID (partition key identifier valid) bit shall be set to zero if the partition key identifier attribute length is zero. Otherwise, the PKI_VALID bit shall be set to one.

The WKI00_VLD (working key identifier 0h valid) bit shall be set to zero if the working key identifier attribute number 8000h has a length of zero. Otherwise, the WKI00_VLD bit shall be set to one.

The WKI01_VLD bit, WKI01_VLD bit, WKI02_VLD bit, WKI03_VLD bit, WKI04_VLD bit, WKI05_VLD bit, WKI06_VLD bit, WKI07_VLD bit, WKI08_VLD bit, WKI09_VLD bit, WKI0A_VLD bit, WKI0B_VLD bit, WKI0C_VLD bit, WKI0D_VLD bit, WKI0E_VLD bit, and WKI0F_VLD have the same bit value definitions as the WKI00_VLD bit, except that the definitions apply to the attributes with numbers 8001h to 800Fh, respectively.

The sixteen WORKING KEY IDENTIFIER fields contain the working key identifier attribute values in ascending attribute number order. If a working key identifier valid bit is set to one, the corresponding WORKING KEY IDENTIFIER field contains the value of the working key identifier attribute. Otherwise, the contents of the WORKING KEY IDENTIFIER field are undefined.

7.1.3.24 Collection Policy/Security attributes page

The Collection Policy/Security attributes page (C+5h) shall contain the attributes listed in table 203.

Table 203 — Collection Policy/Security attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h to 4000 0000h		Reserved	No	
4000 0001h	4	Policy access tag	Yes	Yes
4000 0002h to FFFF FFFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Collection Policy/Security".

The policy access tag attribute (number 4000 0001h) specifies the expected non-zero contents of the POLICY ACCESS TAG field in any capability (see 4.11.2) that allows access to this collection. The format, use, and attribute setting restrictions for the policy access tag attribute are described in 4.11.3.2. A CREATE COLLECTION

command (see 6.8) shall copy the user object policy access tag attribute from the Partition Policy/Security attributes page (see 7.1.3.23) to the policy access tag attribute in new Collection Policy/Security attributes page.

If a command attempts to set an attribute that table 203 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Collection Policy/Security attributes page is shown in table 204.

Table 204 — Collection Policy/Security attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
3	PAGE NUMBER (C+5h)							(LSB)
4	(MSB)							
7	PAGE LENGTH (4h)							(LSB)
8	(MSB)							
11	POLICY ACCESS TAG							(LSB)

The PAGE NUMBER field contains the attributes page number of the Collection Policy/Security attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Collection Policy/Security attributes page.

The POLICY ACCESS TAG field contains the value of the policy access tag attribute.

7.1.3.25 User Object Policy/Security attributes page

The User Object Policy/Security attributes page (5h) shall contain the attributes listed in table 205.

Table 205 — User Object Policy/Security attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h to 4000 0000h		Reserved	No	
4000 0001h	4	Policy access tag	Yes	Yes
4000 0002h to FFFF FFFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 User Object Policy/Security".

The policy access tag attribute (number 4000 0001h) specifies the expected non-zero contents of the POLICY ACCESS TAG field in any capability (see 4.11.2) that allows access to this user object. The format, use, and attribute setting restrictions for the policy access tag attribute are described in 4.11.3.2. A CREATE command (see 6.5) or CREATE AND WRITE command (see 6.6) shall copy the user object policy access tag attribute from the Partition

Policy/Security attributes page (see 7.1.3.23) to the policy access tag attribute in new User Object Policy/Security attributes page.

If a command attempts to set an attribute that table 205 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the User Object Policy/Security attributes page is shown in table 206.

Table 206 — User Object Policy/Security attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
3	PAGE NUMBER (5h)							(LSB)
4	(MSB)							
7	PAGE LENGTH (4h)							(LSB)
8	(MSB)							
11	POLICY ACCESS TAG							(LSB)

The PAGE NUMBER field contains the attributes page number of the User Object Policy/Security attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the User Object Policy/Security attributes page.

The POLICY ACCESS TAG field contains the value of the policy access tag attribute.

7.1.3.26 Root Error Recovery attributes page

The Root Error Recovery attributes page (R+6h) shall contain the attributes listed in table 207.

Table 207 — Root Error Recovery attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	1	Root damage summary	Yes	Yes
2h	1	Contained objects damage summary	No	Yes
3h	6	Last damaged object data time	No	Yes
4h	6	Last damaged object attributes time	No	Yes
5h	6	Last damaged contained object time	No	Yes
6h	8	Number of damaged partitions	No	Yes
7h to FFFF FFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Root Error Recovery".

The root damage summary attribute (1h) indicates the overall error recovery status of the root object using the format shown in table 208.

Table 208 — Root damage summary attribute value

Bit	7	6	5	4	3	2	1	0
	P_OSC	Reserved			P_OSC_RC	R_OSC_RC	ATTR	P_LIST

If the P_OSC (partition object structure check) bit is set to zero, no partitions are processing an OBJECT STRUCTURE CHECK command (see 6.22). If the P_OSC bit is set to one, one or more partitions are processing an OBJECT STRUCTURE CHECK command.

If the P_OSC_RC (partition object structure check recommended) bit is set to zero, the processing of an OBJECT STRUCTURE CHECK command is not recommended for any partitions. If the P_OSC_RC bit is set to one, then one or more partitions may benefit from the processing of an OBJECT STRUCTURE CHECK command. The partition damage summary attribute in each Partition Error Recovery attributes page (see 7.1.3.27) indicates which partitions may benefit from the processing of an OBJECT STRUCTURE CHECK command.

A P_OSC_RC bit that is set to one does not require the processing of an OBJECT STRUCTURE CHECK command on one or more partitions. When the processing of such an OBJECT STRUCTURE CHECK command is required, the process described in 4.11.3.3 is used.

If the R_OSC_RC (root object structure check recommended) bit is set to zero, the processing of an OBJECT STRUCTURE CHECK command is not recommended for the root object and all partitions. If the R_OSC_RC bit is set to one, the processing of an OBJECT STRUCTURE CHECK command is recommended for the root object and all partitions.

An R_OSC_RC bit that is set to one does not require the processing of an OBJECT STRUCTURE CHECK command for the root object and all partitions. When the processing of such an OBJECT STRUCTURE CHECK command is required, the process described in 4.11.3.3 is used.

If the ATTR (attributes) bit is set to zero, no uncorrectable damage has been detected in root object attributes. If the ATTR bit is set to one, uncorrectable damage has been detected in one or more root object attributes.

If the P_LIST (partition list) bit is set to zero, no uncorrectable damage has been detected in the list of partitions in the root object. If the P_LIST bit is set to one, uncorrectable damage has been detected in the list of partitions in the root object.

If the application client sets the root damage summary attribute to any value, the device server shall recompute the attribute's contents.

The contained objects damage summary attribute (2h) indicates the overall error recovery status of all partitions, all collections, and all user objects using the format shown in table 209.

Table 209 — Contained objects damage summary root attribute value

Bit	7	6	5	4	3	2	1	0
	Reserved						C_ATTR	C_DATA

If the C_ATTR (contained attributes) bit is set to zero, no uncorrectable damage has been detected in any attributes associated with a partition, a collection or a user object. If the C_ATTR bit is set to one, uncorrectable damage has been detected in one or more attributes associated with one or more partitions, collections or user objects.

If the C_DATA (contained data) bit is set to zero, no uncorrectable damage has been detected the contained data of any partition, collection, or user object. If the C_DATA bit is set to one, uncorrectable damage has been detected one or more of the following contained data regions:

- a) The list of collections and user objects in one or more partitions;
- b) The list of user objects in one or more collections; or
- c) The user data contained in one or more user objects.

The last damaged object data time attribute (3h) contains the value of the clock attribute in the Root Information attributes page (see 7.1.3.8) when uncorrectable damage was most recently detected in the partition list of the root object. The attribute shall not be modified when an application client corrects the damage. The TIMESTAMPS CONTROL field (see 5.2.13) and the bypass timestamps attribute in the Root Timestamps attributes page (see 7.1.3.15) shall not affect the updating of the last damaged object data time attribute.

The last damaged object attributes time attribute (4h) contains the value of the clock attribute in the Root Information attributes page when uncorrectable damage was most recently detected in a root object attribute. The attribute shall not be modified when an application client corrects the damage. The TIMESTAMPS CONTROL field (see 5.2.13) and the bypass timestamps attribute in the Root Timestamps attributes page shall not affect the updating of the last damaged object attributes time attribute.

The last damaged contained object time attribute (5h) contains the value of the clock attribute in the Root Information attributes page when uncorrectable damage was most recently detected in any of the following:

- a) The list of collections and user objects in one or more partitions;
- b) A partition attribute;
- c) The list of user objects in one or more collections;
- d) A collection attribute;
- e) The user data contained in one or more user objects; or
- f) A user object attribute.

The last damaged contained object time attribute shall not be modified when an application client corrects the damage. The TIMESTAMPS CONTROL field (see 5.2.13) and the bypass timestamps attribute in the Root Timestamps attributes page shall not affect the updating of the last damaged contained object time attribute.

The number of damaged partitions attribute (6h) contains the number of partitions that have unrecovered uncorrectable damage in any of the following:

- a) Their list of member collections and user objects;
- b) A partition attribute;
- c) The list of user objects in one or more member collections;
- d) A collection attribute;
- e) The user data contained in one or more member user objects; or
- f) A user object attribute.

If a command attempts to set an attribute that table 207 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Root Error Recovery attributes page is shown in table 210.

Table 210 — Root Error Recovery attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) PAGE NUMBER (R+6h) (LSB)							
3								
4	(MSB) PAGE LENGTH (1Ch) (LSB)							
7								
8	(MSB) NUMBER OF DAMAGED PARTITIONS (LSB)							
15								
16	ROOT DAMAGE SUMMARY							
17	CONTAINED OBJECTS DAMAGE SUMMARY							
18	(MSB) LAST DAMAGED OBJECT DATA TIME (LSB)							
23								
24	(MSB) LAST DAMAGED OBJECT ATTRIBUTES TIME (LSB)							
29								
30	(MSB) LAST DAMAGED CONTAINED OBJECT TIME (LSB)							
35								

The PAGE NUMBER field contains the attributes page number of the Root Error Recovery attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Root Error Recovery attributes page.

The NUMBER OF DAMAGED PARTITIONS field contains the value of the number of damaged partitions attribute.

The ROOT DAMAGE SUMMARY field contains the value of the root damage summary attribute.

The CONTAINED OBJECTS DAMAGE SUMMARY field contains the value of the contained objects damage summary attribute.

The LAST DAMAGED OBJECT DATA TIME field contains the value of the last damaged object data time attribute.

The LAST DAMAGED OBJECT ATTRIBUTES TIME field contains the value of the last damaged object attributes time attribute.

The LAST DAMAGED CONTAINED OBJECT TIME field contains the value of the last damaged contained object time attribute.

7.1.3.27 Partition Error Recovery attributes page

The Partition Error Recovery attributes page (P+6h) shall contain the attributes listed in table 211.

Table 211 — Partition Error Recovery attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	1	Partition damage summary	Yes	Yes
2h	1	Contained objects damage summary	No	Yes
3h	6	Last damaged object data time	No	Yes
4h	6	Last damaged object attributes time	No	Yes
5h	6	Last damaged contained object time	No	Yes
6h	8	Number of damaged objects	No	Yes
7h to FFFF FFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Partition Error Recovery".

The partition damage summary attribute (1h) indicates the overall error recovery status of the partition using the format shown in table 212.

Table 212 — Partition damage summary attribute value

Bit	7	6	5	4	3	2	1	0
	Reserved				P_OSC_RC	Reserved	ATTR	M_LIST

If the P_OSC_RC (partition object structure check recommended) bit is set to zero, the processing of an OBJECT STRUCTURE CHECK command (see 6.22) is not recommended for the partition. If the P_OSC_RC bit is set to one, the partition may benefit from the processing of an OBJECT STRUCTURE CHECK command.

A P_OSC_RC bit that is set to one does not require the processing of an OBJECT STRUCTURE CHECK command on the partition. When the processing of such an OBJECT STRUCTURE CHECK command is required, the process described in 4.11.3.3 is used.

If the ATTR (attributes) bit is set to zero, no uncorrectable damage has been detected in partition attributes. If the ATTR bit is set to one, uncorrectable damage has been detected in one or more partition attributes.

If the M_LIST (member list) bit is set to zero, no uncorrectable damage has been detected in the list of collections and user objects that are members of the partition. If the M_LIST bit is set to one, uncorrectable damage has been detected in the list of collections and user objects that are members of the partition.

If the application client sets the partition damage summary attribute to any value, the device server shall recompute the attribute's contents.

The contained objects damage summary attribute (2h) indicates the overall error recovery status of all collections and all user objects in the partition using the format shown in table 213.

Table 213 — Contained objects damage summary partition attribute value

Bit	7	6	5	4	3	2	1	0
	Reserved						C_ATTR	C_DATA

If the c_ATTR (contained attributes) bit is set to zero, no uncorrectable damage has been detected in any attributes associated with a collection or a user object in the partition. If the c_ATTR bit is set to one, uncorrectable damage has been detected in one or more attributes associated with one or more collections or user objects in the partition.

If the C_DATA (contained data) bit is set to zero, no uncorrectable damage has been detected the contained data in any collection or user object in the partition. If the C_DATA bit is set to one, uncorrectable damage has been detected one or more of the following contained data regions:

- a) The list of user objects in one or more collections; or
- b) The user data contained in one or more user objects.

The last damaged object data time attribute (3h) contains the value of the clock attribute in the Root Information attributes page (see 7.1.3.8) when uncorrectable damage was most recently detected in the list of collections and user objects that are members the partition. The attribute shall not be modified when an application client corrects the damage. The TIMESTAMPS CONTROL field (see 5.2.13) and the bypass timestamps attribute in the Root Timestamps attributes page (see 7.1.3.15) shall not affect the updating of the last damaged object data time attribute.

The last damaged object attributes time attribute (4h) contains the value of the clock attribute in the Root Information attributes page when uncorrectable damage was most recently detected in a partition attribute. The attribute shall not be modified when an application client corrects the damage. The TIMESTAMPS CONTROL field (see 5.2.13) and the bypass timestamps attribute in the Root Timestamps attributes page shall not affect the updating of the last damaged object attributes time attribute.

The last damaged contained object time attribute (5h) contains the value of the clock attribute in the Root Information attributes page when uncorrectable damage was most recently detected in any of the following:

- a) The list of user objects in one or more member collections;
- b) A collection attribute;
- c) The user data contained in one or more member user objects; or
- d) A user object attribute.

The last damaged contained object time attribute shall not be modified when an application client corrects the damage. The TIMESTAMPS CONTROL field (see 5.2.13) and the bypass timestamps attribute in the Root Timestamps attributes page shall not affect the updating of the last damaged contained object time attribute.

The number of damaged objects attribute (6h) contains the number of member collections and user objects that have unrecovered uncorrectable damage in any of the following:

- a) The list of user objects in one or more member collections;
- b) A collection attribute;
- c) The user data contained in one or more member user objects; or
- d) A user object attribute.

If a command attempts to set an attribute that table 211 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Partition Error Recovery attributes page is shown in table 214.

Table 214 — Partition Error Recovery attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) PAGE NUMBER (P+6h) (LSB)							
3								
4	(MSB) PAGE LENGTH (1Ch) (LSB)							
7								
8	(MSB) NUMBER OF DAMAGED OBJECTS (LSB)							
15								
16	PARTITION DAMAGE SUMMARY							
17	CONTAINED OBJECTS DAMAGE SUMMARY							
18	(MSB) LAST DAMAGED OBJECT DATA TIME (LSB)							
23								
24	(MSB) LAST DAMAGED OBJECT ATTRIBUTES TIME (LSB)							
29								
30	(MSB) LAST DAMAGED CONTAINED OBJECT TIME (LSB)							
35								

The PAGE NUMBER field contains the attributes page number of the Partition Error Recovery attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Partition Error Recovery attributes page.

The NUMBER OF DAMAGED OBJECTS field contains the value of the number of damaged objects attribute.

The PARTITION DAMAGE SUMMARY field contains the value of the partition damage summary attribute.

The CONTAINED OBJECTS DAMAGE SUMMARY field contains the value of the contained objects damage summary attribute.

The LAST DAMAGED OBJECT DATA TIME field contains the value of the last damaged object data time attribute.

The LAST DAMAGED OBJECT ATTRIBUTES TIME field contains the value of the last damaged object attributes time attribute.

The LAST DAMAGED CONTAINED OBJECT TIME field contains the value of the last damaged contained object time attribute.

7.1.3.28 Collection Error Recovery attributes page

The Collection Error Recovery attributes page (C+6h) shall contain the attributes listed in table 215.

Table 215 — Collection Error Recovery attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	1	Collection damage summary	Yes	Yes
2h		Reserved	No	
3h	6	Last damaged object data time	No	Yes
4h	6	Last damaged attributes time	No	Yes
5h to FFFF FFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Collection Error Recovery".

The collection damage summary attribute (1h) indicates the overall error recovery status of the collection using the format shown in table 216.

Table 216 — Collection damage summary attribute value

Bit	7	6	5	4	3	2	1	0
	Reserved						ATTR	C_LIST

If the ATTR (attributes) bit is set to zero, no uncorrectable damage has been detected in collection attributes. If the ATTR bit is set to one, uncorrectable damage has been detected in one or more collection attributes.

If the C_LIST (collection list) bit is set to zero, no uncorrectable damage has been detected in the list of user objects that are members of the collection. If the C_LIST bit is set to one, uncorrectable damage has been detected in the list of user objects that are members of the collection.

If the application client sets the collection damage summary attribute to any value, the device server shall recompute the attribute's contents.

The last damaged object data time attribute (3h) contains the value of the clock attribute in the Root Information attributes page (see 7.1.3.8) when uncorrectable damage was most recently detected in the list of user objects that are members the collection. The attribute shall not be modified when an application client corrects the damage. The TIMESTAMPS CONTROL field (see 5.2.13) and the bypass timestamps attribute in the Root Timestamps attributes page (see 7.1.3.15) shall not affect the updating of the last damaged object data time attribute.

The last damaged object attributes time attribute (4h) contains the value of the clock attribute in the Root Information attributes page when uncorrectable damage was most recently detected in a collection attribute. The attribute shall not be modified when an application client corrects the damage. The TIMESTAMPS CONTROL field (see 5.2.13) and the bypass timestamps attribute in the Root Timestamps attributes page shall not affect the updating of the last damaged object attributes time attribute.

If a command attempts to set an attribute that table 215 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Collection Error Recovery attributes page is shown in table 217.

Table 217 — Collection Error Recovery attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
3	PAGE NUMBER (C+6h)							(LSB)
4	(MSB)							
7	PAGE LENGTH (Eh)							(LSB)
8	COLLECTION DAMAGE SUMMARY							
9	Reserved							
10	(MSB)							
15	LAST DAMAGED OBJECT DATA TIME							(LSB)
16	(MSB)							
21	LAST DAMAGED ATTRIBUTES TIME							(LSB)

The PAGE NUMBER field contains the attributes page number of the Collection Error Recovery attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Collection Error Recovery attributes page.

The COLLECTION DAMAGE SUMMARY field contains the value of the collection damage summary attribute.

The LAST DAMAGED OBJECT DATA TIME field contains the value of the last damaged object data time attribute.

The LAST DAMAGED ATTRIBUTES TIME field contains the value of the last damaged attributes time attribute.

7.1.3.29 User Object Error Recovery attributes page

The User Object Error Recovery attributes page (6h) shall contain the attributes listed in table 218.

Table 218 — User Object Error Recovery attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	1	User object damage summary	Yes	Yes
2h		Reserved	No	
3h	6	Last damaged object data time	No	Yes
4h	6	Last damaged attributes time	No	Yes
5h to FFFF FFFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the **VENDOR IDENTIFICATION** field containing the ASCII characters "INCITS" and the **ATTRIBUTES PAGE IDENTIFICATION** field containing the ASCII characters "T10 User Object Error Recovery".

The user object damage summary attribute (1h) indicates the overall error recovery status of the user object using the format shown in table 219.

Table 219 — User object damage summary attribute value

Bit	7	6	5	4	3	2	1	0
	Reserved						ATTR	DATA

If the **ATTR** (attributes) bit is set to zero, no uncorrectable damage has been detected in user object attributes. If the **ATTR** bit is set to one, uncorrectable damage has been detected in one or more user object attributes.

If the **DATA** bit is set to zero, no uncorrectable damage has been detected in user object's user data. If the **DATA** bit is set to one, uncorrectable damage has been detected in user object's user data. The **READ MAP** command (see 6.28) may be used to determine details of the uncorrectable damage in a user object's user data.

If the application client sets the user object damage summary attribute to any value, the device server shall recompute the attribute's contents.

The last damaged object data time attribute (3h) contains the value of the clock attribute in the Root Information attributes page (see 7.1.3.8) when uncorrectable damage was most recently detected in user object's user data. The attribute shall not be modified when an application client corrects the damage. The **TIMESTAMPS CONTROL** field (see 5.2.13) and the bypass timestamps attribute in the Root Timestamps attributes page (see 7.1.3.15) shall not affect the updating of the last damaged object data time attribute.

The last damaged object attributes time attribute (4h) contains the value of the clock attribute in the Root Information attributes page when uncorrectable damage was most recently detected in a user object attribute. The attribute shall not be modified when an application client corrects the damage. The **TIMESTAMPS CONTROL** field (see 5.2.13) and the bypass timestamps attribute in the Root Timestamps attributes page shall not affect the updating of the last damaged object attributes time attribute.

If a command attempts to set an attribute that table 218 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the User Object Error Recovery attributes page is shown in table 220.

Table 220 — User Object Error Recovery attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	PAGE NUMBER (6h)						(LSB)
3								
4	(MSB)	PAGE LENGTH (Eh)						(LSB)
7								
8		USER OBJECT DAMAGE SUMMARY						
9		Reserved						
10	(MSB)	LAST DAMAGED OBJECT DATA TIME						(LSB)
15								
16	(MSB)	LAST DAMAGED ATTRIBUTES TIME						(LSB)
21								

The PAGE NUMBER field contains the attributes page number of the Collection Error Recovery attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Collection Error Recovery attributes page.

The USER OBJECT DAMAGE SUMMARY field contains the value of the user object damage summary attribute.

The LAST DAMAGED OBJECT DATA TIME field contains the value of the last damaged object data time attribute.

The LAST DAMAGED ATTRIBUTES TIME field contains the value of the last damaged attributes time attribute.

7.1.3.30 Snapshots Information attributes page

The Snapshots Information attributes page (P+7h) shall contain the attributes listed in table 221.

Table 221 — Snapshots Information attributes page contents

Attribute Number	Length (bytes) ^a	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	0 or 40	Page identification	No	Yes
1h	0 or 1	Partition type	No	Yes
2h to 7Fh		Reserved	No	
80h	0 or 8	Source partition	No	Yes
81h	0 or 8	Snapshot backward	No	Yes
82h	0 or 8	Snapshot forward	No	Yes
83h to FFFFh	0 or 8	Clone destination	No	Yes
1 0000h to 2 0000h		Reserved	No	
2 0001h	0 or 4	Snapshots count	No	Yes
2 0002h	0 or 4	Clones count	No	Yes
2 0003h to 2 000Bh		Reserved	No	
2 000Ch	0 or 4	Branch depth	No	Yes
2 000Dh to 2 0010h		Reserved	No	
2 0011h	0 or 6	Create completion time	No	Yes
2 0012h	0 or 6	Refresh completion time	No	Yes
2 0013h	0 or 6	Restore completion time	No	Yes
2 0014h	0 or 8	Restore Partition_ID	No	Yes
2 0015 to FFFF FFEh		Reserved	No	
^a A length of 0 in this column denotes an attribute that may be undefined (see 3.1.51).				

If it is defined (see 3.1.14), the page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Snapshots Information".

If it is defined (see 3.1.14), the partition type attribute (number 1h) (see table 222) indicates the characteristics of the partition with respect to the snapshots model (see 4.13.2). If the partition type attribute is undefined (see 3.1.51), the partition is a primary partition.

Table 222 — Partition type attribute values

Partition type attribute value	Description
00h	Primary partition (i.e., not a snapshot partition or a clone partition)
01h	Snapshot partition
02h	Clone partition
03h to FFh	Reserved

If it is defined (see 3.1.14), the source partition attribute (number 80h) contains the contents of the SOURCE PARTITION_ID field in the CREATE SNAPSHOT command (see 6.10) or CREATE CLONE command (see 6.7) that created the partition. If the source partition attribute is undefined (see 3.1.14), then one of the following is true:

- a) The partition was not created by a CREATE SNAPSHOT command or a CREATE CLONE command, or
- b) The partition was created by a CREATE CLONE command and later detached by a DETACH CLONE command (see 6.12).

If it is defined (see 3.1.14), the snapshot backward attribute (number 81h) contains the Partition_ID (see 4.6.2) of the next older snapshot partition in the history chain (see 4.13.2.2). If the snapshot backward attribute is undefined (see 3.1.51), then the partition has never been a source partition or a destination partition in a CREATE SNAPSHOT command (see 6.10).

If it is defined (see 3.1.14), the snapshot forward attribute (number 82h) contains the Partition_ID (see 4.6.2) of the next newer snapshot partition in the history chain (see 4.13.2.2). If the snapshot forward attribute is undefined (see 3.1.51), then the partition has never been a source partition or a destination partition in a CREATE SNAPSHOT command (see 6.10).

Each defined (see 3.1.14) clone destination attribute (numbers 83h to FFFFh) contains the Partition_ID (see 4.6.2) of a clone partition (see 4.13.2.3). If all clone destination attributes are undefined (see 3.1.51), then one of the following is true:

- a) The partition has never been a source partition for a CREATE CLONE command (see 6.12), or
- b) All clone partitions for which this partition was the source have been:
 - A) Detached by DETACH CLONE commands (see 6.12); or
 - B) Removed by REMOVE PARTITION commands (see 6.34).

There is no significance to which clone destination attribute numbers are defined and which are undefined.

If the snapshot backward attribute is defined (see 3.1.14) in a primary partition or clone partition, then the snapshots count attribute (number 2 001h) is defined and contains the number of snapshots in the history chain (see 4.13.2.2) that the primary partition or clone partition heads. If the snapshot backward attribute is undefined (see 3.1.51) or the partition is a snapshot partition, the snapshots count attribute is undefined.

If any clone destination attribute is defined (see 3.1.14), then the clones count attribute (number 2 002h) is defined and contains the number of clone destination attributes that are defined in the partition. If all clone destination attributes are undefined (see 3.1.51), the clones count attribute is undefined.

If it is defined (see 3.1.14), the create completion time attribute (number 2 0011h) contains value of the clock attribute in the Root Information attributes page (see 7.1.3.8) at the completion of the CREATE SNAPSHOT command (see 6.10) or CREATE CLONE command (see 6.7) that created this partition. The create completion time attribute is undefined (see 3.1.51) if any of the following are true:

- a) The partition was not created by a CREATE SNAPSHOT command or a CREATE CLONE command; or
- b) The CREATE SNAPSHOT command or CREATE CLONE command has not yet completed.

If it is defined (see 3.1.14), the refresh completion time attribute (number 2 0012h) contains value of the clock attribute in the Root Information attributes page (see 7.1.3.8) at the completion of the most recent REFRESH SNAPSHOT command (see 6.30). The refresh completion time attribute is undefined (see 3.1.51) if any of the following are true:

- a) The partition has never been the destination of a REFRESH SNAPSHOT command; or
- b) The most recent REFRESH SNAPSHOT command has not yet completed.

If it is defined (see 3.1.14), the restore completion time attribute (number 2 0013h) contains value of the clock attribute in the Root Information attributes page (see 7.1.3.8) at the completion of the most recent RESTORE PARTITION FROM SNAPSHOT command (see 6.35). The restore completion time attribute is undefined (see 3.1.51) if any of the following are true:

- a) The partition has never been the destination of a RESTORE PARTITION FROM SNAPSHOT command; or
- b) The most recent RESTORE PARTITION FROM SNAPSHOT command has not yet completed.

If it is defined (see 3.1.14), the restore Partition_ID attribute (number 2 0014h) contains value in the PARTITION_ID field of the most recent RESTORE PARTITION FROM SNAPSHOT command (see 6.35) that has completed. The restore Partition_ID attribute is undefined (see 3.1.51) if any of the following are true:

- a) The partition has never been the destination of a RESTORE PARTITION FROM SNAPSHOT command; or
- b) The most recent RESTORE PARTITION FROM SNAPSHOT command has not yet completed.

If it is defined (see 3.1.14), the branch depth attribute (number 0002 000Ch) indicates the nesting depth of a snapshot partition or clone partition. The branch depth of a primary partition is zero, and the branch depth attribute is undefined (see 3.1.51) for primary partitions. Other branch depth values increase from the primary partition value as follows:

- a) All snapshot partitions that have the primary partition as their source partition have a branch depth of zero;
- b) All clone partitions that have a snapshot partition with a branch depth of zero as their source partition, have a branch depth of one;
- c) All snapshot partitions that have a clone partition with a branch depth of one as their source partition, have a branch depth of one;
- d) All clone partitions that have a snapshot with a branch depth of one as their source partition, have a branch depth of two;
- e) All snapshot partitions that have a clone partition with a branch depth of n as their source partition, have a branch depth of n; and
- f) All clone partitions that have a snapshot partition with a branch depth of n as their source partition, have a branch depth of n plus one.

If a command attempts to set an attribute that table 221 states is not application client settable, then the command shall be terminated as described in 7.1.2.

7.1.3.31 Current Command attributes page

The Current Command attributes page (FFFF FFFEh) shall contain the attributes listed in table 223.

Table 223 — Current Command attributes page contents

Attribute Number	Length (bytes)	Attribute	Application Client Settable	OSD Logical Unit Provided
0h	40	Page identification	No	Yes
1h	32	Response integrity check value	No	Yes
2h	1	Object Type	No	Yes
3h	8	Partition_ID	No	Yes
4h	8	Collection_Object_ID or User_Object_ID	No	Yes
5h	8	Starting byte address of append	No	Yes
6h to FFFF FFFEh		Reserved	No	

The page identification attribute (number 0h) shall have the format described in 7.1.3.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Current Command".

If the NOSEC security method or the CAPKEY security method (see 4.12.4) is used to process the command or if status returned for the command is CHECK CONDITION status, the response integrity check value attribute (number 1h) shall contain zero. Otherwise, the response integrity check value attribute shall contain an integrity check value (see 4.12.8) that is computed as described in 4.12.4.4.

NOTE 10 If a command terminates with CHECK CONDITION status, the response integrity check value is returned in the sense data (see 4.16).

The object type attribute (number 2h) shall identify the type of OSD object on which the current command is operating using the code values shown in table 17 (see 4.11.2.2).

The Partition_ID attribute (number 3h) shall contain the Partition_ID (see 4.6.4) of partition containing the OSD object on which the current command is operating.

If the object type attribute contains COLLECTION (see table 17 in 4.11.2.2), the Collection_Object_ID or User_Object_ID attribute (number 4h) shall contain the Collection_Object_ID (see 4.6.6) of the collection on which the current command is operating. Otherwise, the Collection_Object_ID or User_Object_ID attribute shall contain the User_Object_ID (see 4.6.5) of the user object on which the current command is operating.

If the current command is an APPEND (see 6.2), the starting byte address of append attribute (number 5h) shall contain the starting byte address used for the append command function. If the current command is not an APPEND, the starting byte address of append attribute shall contain zero.

If a command attempts to set an attribute that table 223 states is not application client settable, then the command shall be terminated as described in 7.1.2.

The page format for the Current Command attributes page is shown in table 224.

Table 224 — Current Command attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	PAGE NUMBER (FFFF FFEh)						(LSB)
3								
4	(MSB)	PAGE LENGTH (3Ch)						(LSB)
7								
8	(MSB)	RESPONSE INTEGRITY CHECK VALUE						(LSB)
39								
40		OBJECT TYPE						
41								
43		Reserved						
44	(MSB)							(LSB)
51		PARTITION_ID						
52	(MSB)							(LSB)
59		COLLECTION_OBJECT_ID OR USER_OBJECT_ID						
60	(MSB)							(LSB)
67		STARTING BYTE ADDRESS OF APPEND						
								(LSB)

The PAGE NUMBER field contains the attributes page number of the Current Command attributes page.

The PAGE LENGTH field contains the number of additional bytes in the page format of the Current Command attributes page.

The RESPONSE INTEGRITY CHECK VALUE field contains the value of the response integrity check value attribute.

The OBJECT TYPE field contains the value of the object type attribute.

The PARTITION_ID field contains the value of the Partition_ID attribute.

The COLLECTION_OBJECT_ID OR USER_OBJECT_ID field contains the value of the Collection_Object_ID or User_Object_ID attribute.

The STARTING BYTE ADDRESS OF APPEND field contains the value of the starting byte address of append attribute.

7.1.3.32 Null attributes page

The page format for the null attributes page is shown in table 225.

Table 225 — Null attributes page format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) PAGE NUMBER							(LSB)
3								
4	(MSB) PAGE LENGTH (00h)							(LSB)
7								

The PAGE NUMBER field contains the attributes page number of the requested attributes page.

The PAGE LENGTH field contains zero.

7.1.4 OSD attributes lists

7.1.4.1 Attributes lists overview

An attributes list acts on one or more individual attribute values using attributes page and attribute number values to specify the attribute values to be retrieved or set.

The format of an attributes list is shown in table 226.

Table 226 — Attributes list format

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved				LIST TYPE			
1	Reserved							
3								
4	(MSB) LIST LENGTH (n-7)							(LSB)
7								
	Attributes list entries							
8	Attributes list entry 0							
	⋮							
n	Attributes list entry x							

The LIST TYPE field (see table 227) specifies the format of all attributes list entries in the attributes list.

Table 227 — List type values

List Type	Description	Support Requirement	Reference	Allowed Use		
				Get Attributes		Set Attributes List
				List	Response	
0h	Reserved			No	No	No
1h	Retrieve attributes for the specified OSD object or objects	Mandatory	7.1.4.2	Yes	No	No
2h to 8h	Reserved			No	No	No
9h	Retrieved attributes for the specified OSD object, or set attributes for the specified OSD object or objects	Mandatory	7.1.4.3	No	Yes	Yes
Ah to Dh	Reserved			No	No	No
Eh	Retrieved attributes for more than one OSD object	Mandatory	7.1.4.4	No	Yes	No
Fh	Obsolete					

If list type 1h (see 7.1.4.2) is used to retrieve attributes for the specified OSD object, the list type of the list containing the retrieved objects shall be:

- a) Eh (see 7.1.4.4) for:
 - A) A CREATE command that creates more than one user object; or
 - B) A GET MEMBER ATTRIBUTES command; or
 - C) A SET MEMBER ATTRIBUTES command;
 or
- b) 9h (see 7.1.4.3) for all other commands and for a CREATE command that creates only one user object.

The LIST LENGTH field indicates the number of bytes of attributes list entries that follow or the number of bytes that would follow if the parameter data were not truncated (see 5.2.6.4). For attributes lists sent from the application client to the device server, the LIST LENGTH field may contain zero.

For an attributes list sent from the device server to the application client, a list length of zero indicates that all of the requested attributes are undefined attributes (see 3.1.51).

The application client should set the list length to zero in any attributes list that it sends to the device server. The device server shall use the length of the list specified in the CDB and shall ignore the contents of the LIST LENGTH field.

7.1.4.2 List entry format for retrieving attributes for a specified OSD object

The attributes list entry format shown in table 228 is used for specifying the attributes to be retrieved by a GET ATTRIBUTES command (see 6.18) or equivalent command function.

For the GET MEMBER ATTRIBUTES command (see 6.19) and the SET MEMBER ATTRIBUTES command (see 6.39), the list entry format shown in table 228 specifies retrieval of the same attributes (i.e., the same combinations of attribute page and attribute number) for all user objects in the specified collection.

Table 228 — List entry format for retrieving attributes for a specified OSD object

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
3	ATTRIBUTES PAGE							(LSB)
4	(MSB) _____							
7	ATTRIBUTE NUMBER							(LSB)

The ATTRIBUTES PAGE field specifies the page number of one attribute to be returned. If the specified attributes page number is not associated with an object specified by the command, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

NOTE 11 Some commands (e.g., LIST (see 6.20)) define methods that allow the attributes in multiple objects to be processed by a single command. The conditions under which an attributes page number is not required to be associated with the object specified by a command's CDB appear in the command definition subclauses for the exceptional commands.

The ATTRIBUTE NUMBER field specifies:

- The attribute number within the attributes page specified by the ATTRIBUTES PAGE field of the one attribute value to be returned; or
- The value FFFF FFFFh to request the return of each defined attribute (see 3.1.14) in the attributes page specified by the ATTRIBUTES PAGE field.

Requirements on the ATTRIBUTE LENGTH field for retrieved attributes are described in 4.8.2.

Specifying attributes page and attribute number values of FFFF FFFFh causes all defined attributes values in all defined pages associated with the OSD object specified by a command to be returned. Specifying an attribute numbers value of FFFF FFFFh causes all defined attributes values in the specified attributes page to be returned.

If FFFF FFFFh is used as an attributes page number or attribute number value, only defined attributes (see 3.1.14) shall be returned.

7.1.4.3 List entry format for retrieved attributes and for setting attributes for the specified OSD object

The attributes list entry format shown in table 229 is used for returning the each attribute value to be retrieved by a GET ATTRIBUTES command (see 6.18) and for specifying each attribute value to be set by a SET ATTRIBUTES command (see 6.36) or equivalent command functions.

For the GET MEMBER ATTRIBUTES command (see 6.19) and the SET MEMBER ATTRIBUTES command (see 6.39), the list entry format shown in table 229 specifies the same attributes values to set in all user objects in the specified collection. The retrieved attributes format for the GET MEMBER ATTRIBUTES command and SET MEMBER ATTRIBUTES command is described in 7.1.4.4.

Table 229 — List entry format for retrieved attributes and for setting attributes for the specified OSD object

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	ATTRIBUTES PAGE						(LSB)
3								
4	(MSB)	ATTRIBUTE NUMBER						(LSB)
7								
8	(MSB)	ATTRIBUTE LENGTH (n-9)						(LSB)
9								
10	(MSB)	ATTRIBUTE VALUE						(LSB)
n								
n+1		Pad bytes (for eight-byte alignment)						
m								

The ATTRIBUTES PAGE field specifies the page number of the attribute value.

The ATTRIBUTE NUMBER field specifies the attribute number within the attributes page specified by the ATTRIBUTES PAGE field of the attribute value.

The ATTRIBUTE LENGTH field specifies the length of the attribute value in bytes. The contents of the ATTRIBUTE LENGTH field and ATTRIBUTE VALUE field for retrieved attributes are described in 4.8.2. The effects of the ATTRIBUTE LENGTH field on the setting of attributes are described in 4.8.3.

The ATTRIBUTE VALUE field specifies the attribute value.

If the ATTRIBUTES PAGE or ATTRIBUTE NUMBER field contains FFFF FFFFh for a set command function, the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The list entry length shall be a multiple of eight bytes. Depending on the attribute length, zero to seven bytes containing zeros shall be added at the end of the list entry to meet the length requirement.

7.1.4.4 Multi-object retrieved attributes format

The format shown in table 230 is used for indicating the attributes to be retrieved by:

- a) A CREATE command (see 6.5) that creates more than one user object;
- b) A GET MEMBER ATTRIBUTES command (see 6.19);
- c) SET MEMBER ATTRIBUTES command (see 6.39);
- d) A LIST command (see 6.20) with the LIST_ATTR bit set to one; and
- e) A LIST COLLECTION command (see 6.21) with the LIST_ATTR bit set to one.

Table 230 — Multi-object retrieved attributes list format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
7	OBJECT ID (LSB)							
8	OBJECT TYPE							
9	Reserved							
13								
14	(MSB)							
15	ATTRIBUTES LIST LENGTH (n-15) (LSB)							
	Attributes list entries							
16	Attributes list entry 0 (see 7.1.4.3)							
	⋮							
n	Attributes list entry x (see 7.1.4.3)							

The contents of the OBJECT ID field depend on the command that is retrieving the attributes as follows:

- a) For a CREATE command, GET MEMBER ATTRIBUTES command, or SET MEMBER ATTRIBUTES command, the OBJECT ID field contains a User_Object_ID (see 4.6.5);
- b) For a LIST command with the LIST_ATTR bit set to one, the OBJECT ID field contains a Partition_ID (see 4.6.4) or a User_Object_ID as described in 6.20.2; and
- c) For a LIST COLLECTION command with the LIST_ATTR bit set to one, the OBJECT ID field contains a Collection_Object_ID (see 4.6.6) or a User_Object_ID as described in 6.21.

The OBJECT TYPE field indicates type of OSD object to which the attributes list entry applies using the code values shown in table 17 (see 4.11.2.2).

The ATTRIBUTES LIST LENGTH field indicates the number of bytes of attributes list entries that follow. If the parameter data is truncated due to insufficient allocation length, the ATTRIBUTES LIST LENGTH field shall not be altered to reflect the truncation (i.e., the attributes list length indicates the number of bytes that would follow if the allocation length had been infinite).

Each attributes list entry has the format shown in 7.1.4.3 and contains information about one attribute.

7.2 Diagnostic parameters

This subclause defines the descriptors and pages for diagnostic parameters used with OSD type devices.

The diagnostic parameter list is described in SPC-3.

See SPC-3 for diagnostic pages used with all device types.

No diagnostic pages are defined for specific use by OSD type devices.

7.3 Log parameters

This subclause defines the descriptors and pages for log parameters used with OSD type devices.

The log parameter list is described in SPC-3.

See SPC-3 for log parameter pages used with all device types.

No log parameter pages are defined for specific use by OSD type devices.

7.4 Mode parameters

This subclause defines the descriptors and pages for mode parameters used with OSD type devices.

The mode parameter list, including the mode parameter header and mode block descriptor, are described in SPC-3.

OSD type devices shall reserve the following mode parameter header fields (see SPC-3):

- a) MEDIUM TYPE;
- b) DEVICE-SPECIFIC PARAMETER; and
- c) LONGLBA.

OSD type devices shall set the BLOCK DESCRIPTOR LENGTH field to zero and shall return CHECK CONDITION status for any command that attempts to set the block descriptor length to a value other than zero.

See SPC-3 for mode pages used with all device types.

No mode pages are defined for specific use by OSD type devices.

7.5 Vital product data parameters

7.5.1 Overview

This subclause defines the VPD pages used with OSD type devices.

See SPC-3 for VPD pages used with all device types.

The VPD page codes that are specific to OSD type devices are shown in table 231.

Table 231 — OSD specific VPD page codes

Page code	Description	Reference	Support Requirements
B0h	OSD Information	7.5.2	Optional
B1h	Security Token	7.5.3	Optional
B2h to BFh	Reserved for OSD type devices		

7.5.2 OSD Information VPD page

7.5.2.1 Overview

The OSD Information VPD page (see table 232) contains information about the OSD logical unit that may be needed to properly prepare OSD commands.

Table 232 — OSD Information VPD page format

Bit Byte	7	6	5	4	3	2	1	0
0	PERIPHERAL QUALIFIER			PERIPHERAL DEVICE TYPE				
1	PAGE CODE (B0h)							
2	PAGE LENGTH (n-3)							
3								
	OSD information descriptors							
4	OSD information descriptor [first]							
	⋮							
n	OSD information descriptor [last]							

The PERIPHERAL QUALIFIER field and the PERIPHERAL DEVICE TYPE field are defined in SPC-3.

The PAGE LENGTH field specifies the length of the following VPD page data. If the allocation length is less than the length of the data to be returned, the page length shall not be adjusted to reflect the truncation.

Each OSD information descriptor (see table 233) contains information about the OSD logical unit that may be needed to properly prepare OSD commands.

Table 233 — OSD information descriptor format

Bit Byte	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE							
1	Reserved							
2	ADDITIONAL LENGTH (n-3)							
3								
4	Descriptor-specific information							
n								

The DESCRIPTOR TYPE field (see table 234) indicates the format of and information in the OSD information descriptor.

Table 234 — OSD information descriptor type values

Value		Reference	Support Requirements
00h	OSD logical unit security methods	7.5.2.2	Optional
01h to F0h	Reserved		
F1 to FFh	Vendor specific		

The ADDITIONAL LENGTH field specifies the length of the following OSD information descriptor data. If the allocation length causes an OSD information descriptor to be truncated, the additional length shall not be adjusted to reflect the truncation.

The format and content of the descriptor-specific information depends on the descriptor type.

7.5.2.2 OSD logical unit security methods information descriptor

Each OSD logical unit security methods information descriptor (see table 235) contains information about the OSD logical unit security methods that may need to be obtained in order to properly prepare OSD commands.

Table 235 — OSD logical unit security methods information descriptor format

Bit Byte	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE (00h)							
1	Reserved							
2	ADDITIONAL LENGTH (0002h)							
3								
4	ROOT SECURITY METHOD							
5	PARTITION ZERO SECURITY METHOD							

The DESCRIPTOR TYPE field set to 00h indicates that this is an OSD logical unit security methods information descriptor.

The ADDITIONAL LENGTH field specifies the length of the following OSD information descriptor data.

The ROOT SECURITY METHOD field contains the value in the security method attribute in the Root Policy/Security attributes page (see 7.1.3.22).

The PARTITION ZERO SECURITY METHOD field contains the value in the security method attribute in the Partition Policy/Security attributes page (see 7.1.3.23) associated with partition zero.

7.5.3 Security Token VPD page

The Security Token VPD page (see table 236) contains a security token for use in the CAPKEY security method (see 4.12.4.3).

Table 236 — Security Token VPD page

Bit Byte	7	6	5	4	3	2	1	0
0	PERIPHERAL QUALIFIER			PERIPHERAL DEVICE TYPE				
1	PAGE CODE (B1h)							
2	PAGE LENGTH (n-3)							
3								
4	SECURITY TOKEN							
n								

The PERIPHERAL QUALIFIER field and the PERIPHERAL DEVICE TYPE field are defined in SPC-3.

The PAGE LENGTH field indicates the length of the following VPD page data. The page length shall be at least sixteen. If the allocation length is less than the length of the data to be returned, the page length shall not be adjusted to reflect the truncation.

The SECURITY TOKEN field contains a value that is unique to the I_T_L nexus that sent the INQUIRY command. The security token shall be random as defined by RFC 1750. An I_T nexus loss event, logical unit reset event, or reset event (see SAM-4) shall cause the security token to change.

Annex A

(Normative)

Attributes page numbers assigned by other standards

A.1 Attributes page numbers assigned by other standards

At the time of publication, no attribute page numbers are assigned by other standards. The attributes page numbers available for assignment by other standards are shown in table A.1.

Table A.1 — Attributes page numbers assigned by other standards

Page Number	Associated object type	Assignment
R+8000h to R+FFFFh	Root	Reserved
P+8000h to P+FFFFh	Partition	Reserved
C+8000h to C+FFFFh	Collection	Reserved
8000h to EFFFh	User Object	Reserved

Annex B

(Informative)

Numeric order codes

B.1 Service action codes

The variable length CDB service action codes assigned by this standard are shown in table B.1.

Table B.1 — Numerical order OSD service action codes

Service Action	Command
8801h to 8803h	Obsolete
8804h	Reserved
8805h to 8808h	Obsolete
8809h	Reserved
880Ah to 880Ch	Obsolete
880Dh	Reserved
880Eh to 880Fh	Obsolete
8810h to 8811h	Reserved
8812h	Obsolete
8813h to 8814h	Reserved
8815h to 881Ch	Obsolete
881Dh to 887fh	Reserved
8880h	OBJECT STRUCTURE CHECK
8881h	FORMAT OSD
8882h	CREATE
8883h	LIST
8884h	PUNCH
8885h	READ
8886h	WRITE
8887h	APPEND
8888h	FLUSH
8889h	CLEAR
888Ah	REMOVE
888Bh	CREATE PARTITION
888Ch	REMOVE PARTITION
888Dh	Reserved
888Eh	GET ATTRIBUTES
888Fh	SET ATTRIBUTES
8890h to 8891h	Reserved
8892h	CREATE AND WRITE
8893h	COPY USER OBJECTS
8894h	CREATE TRACKING COLLECTION

Table B.1 — Numerical order OSD service action codes

Service Action	Command
8895h	CREATE COLLECTION
8896h	REMOVE COLLECTION
8897h	LIST COLLECTION
8898h	SET KEY
8899h	SET MASTER KEY
889Ah	FLUSH COLLECTION
889Bh	FLUSH PARTITION
889Ch	FLUSH OSD
889Dh to 889Fh	Reserved
88A0h	QUERY
88A1h	REMOVE MEMBER OBJECTS
88A2h	GET MEMBER ATTRIBUTES
88A3h	SET MEMBER ATTRIBUTES
88A4h to 88A7h	Reserved
88A8h	CREATE CLONE
88A9h	CREATE SNAPSHOT
88AAh	DETACH CLONE
88ABh	REFRESH SNAPSHOT OR CLONE
88ACh	RESTORE PARTITION FROM SNAPSHOT
88ADh to 88B0h	Reserved
88B1h	READ MAP
88B2h	READ MAPS AND COMPARE
88B3h to 8F7Bh	Reserved
8F7Ch	PERFORM SCSI COMMAND
8F7Dh	PERFORM TASK MANAGEMENT FUNCTION
8F7Eh to 8F7Fh	Obsolete
8F80h to 8FFFh	Vendor specific

Annex C

(Informative)

Attributes defined by this standard

C.1 Attributes list

The attributes defined by this standard are shown in table C.1.

Table C.1 — Numerical order attributes defined by this standard (part 1 of 6)

Page Number	Page Name	Attribute Number	Attribute
0h	User Object Directory	0h	"INCITS T10 User Object Directory"
		1h	"INCITS T10 User Object Information"
		2h	"INCITS T10 User Object Quotas"
		3h	"INCITS T10 User Object Timestamps"
		4h	"INCITS T10 Collections"
		5h	"INCITS T10 User Object Policy/Security"
		6h	"INCITS T10 User Object Error Recovery"
1h	User Object Information	0h	Page identification
		1h	Partition_ID
		2h	User_Object_ID
		9h	Username
		81h	Used capacity
		82h	User object logical length
		83h	Object accessibility
		D1h	Actual data space
		D2h	Reserved data space
2h	User Object Quotas	0h	Page identification
		1h	Maximum user object length
3h	User Object Timestamps	0h	Page identification
		1h	Created time
		2h	Attributes accessed time
		3h	Attributes modified time
		4h	Data accessed time
		5h	Data modified time
4h	Collections	0h	Page identification
		1h	Collection pointer
		...	" "
		FFFF FF00h	" "
5h	User Object Policy/Security	0h	Page identification
		4000 0001h	Policy access tag

Table C.1 — Numerical order attributes defined by this standard (part 2 of 6)

Page Number	Page Name	Attribute Number	Attribute
6h	User Object Error Recovery	0h	Page identification
		1h	User object damage summary
		3h	Last damaged object data time
		4h	Last damaged attributes time
3000 0000h	Partition Directory	3000 0000h	"INCITS T10 Partition Directory"
		3000 0001h	"INCITS T10 Partition Information"
		3000 0002h	"INCITS T10 Partition Quotas"
		3000 0003h	"INCITS T10 Partition Timestamps"
		3000 0004h	"INCITS T10 Attributes Access"
		3000 0005h	"INCITS T10 Partition Policy/Security"
		3000 0006h	"INCITS T10 Partition Error Recovery"
3000 0001h	Partition Information	3000 0007h	"INCITS T10 Snapshots Information"
		0h	Page identification
		1h	Partition_ID
		9h	Username
		81h	Used capacity
		83h	Object accessibility
		84h	Potential used capacity increment
		C1h	Number of collections and user objects
		D1h	Actual data space
		D2h	Reserved data space
		200h	Default snapshot duplication method
		201h	Default clone duplication method
		202h	Default copy user objects duplication method
		300h	Default snapshot time of duplication method
		301h	Default clone time of duplication method
		302h	Default copy user objects time of duplication method
3000 0002h	Partition Quotas	0h	Page identification
		1h	Default maximum user object length
		1 0001h	Capacity quota
		1 0002h	Object count
		1 0081h	Collections per user object
3000 0003h	Partition Timestamps	0h	Page identification
		1h	Created time
		2h	Attributes accessed time
		3h	Attributes modified time
		4h	Data accessed time
		5h	Data modified time
		FFFF FFEh	Timestamp bypass

Table C.1 — Numerical order attributes defined by this standard (part 3 of 6)

Page Number	Page Name	Attribute Number	Attribute
3000 0004h	Attributes Access	0h	Page identification
		1h	Allowed attributes access
		...	" " "
		FFFF FFEh	" " "
3000 0005h	Partition Policy/Security	0h	Page identification
		1h	Default security method
		2h	Oldest valid nonce
		3h	Newest valid nonce
		4h	Request nonce list depth
		5h	Frozen working key bit mask
		7FFFh	Partition key identifier
		8000h	Working key identifier
		...	" " "
		800Fh	" " "
		4000 0001h	Policy access tag
		4000 0002h	User object policy access tag
3000 0006h	Partition Error Recovery	0h	Page identification
		1h	Partition damage summary
		2h	Contained objects damage summary
		3h	Last damaged object data time
		4h	Last damaged object attributes time
		5h	Last damaged contained object time
		6h	Number of damaged objects
3000 0007h	Snapshots Information	0h	Page identification
		1h	Partition type
		80h	Source partition
		81h	Snapshot backward
		82h	Snapshot forward
		83h	Clone destination
		...	" "
		FFFFh	" "
		2 0001h	Snapshots count
		2 0002h	Clones count
		2 000Ch	Branch depth
		2 0011h	Create completion time
		2 0012h	Refresh completion time
		2 0013h	Restore completion time
		2 0014h	Restore Partition_ID

Table C.1 — Numerical order attributes defined by this standard (part 4 of 6)

Page Number	Page Name	Attribute Number	Attribute
6000 0000h	Collection Directory	6000 0000h 6000 0001h 6000 0003h 6000 0004h 6000 0005h 6000 0006h	"INCITS T10 Collection Directory" "INCITS T10 Collection Information" "INCITS T10 Collection Timestamps" "INCITS T10 Command Tracking" "INCITS T10 Collection Policy/Security" "INCITS T10 Collection Error Recovery"
6000 0001h	Collection Information	0h 1h 2h 9h Ah 81h 83h	Page identification Partition_ID Collection_Object_ID Username Collection type Used capacity Object accessibility
6000 0003h	Collection Timestamps	0h 1h 2h 3h 4h 5h	Page identification Created time Attributes accessed time Attributes modified time Data accessed time Data modified time
6000 0004h	Command Tracking	0h 1h 2h 3h 4h 10h 11h 12h 13h F000 0000h ... FFFF FFFEh	Page identification Percent complete Active command status Ended command status Sense data Number of members Objects processed Newer objects skipped Missing objects skipped Vendor specific " " " "
6000 0005h	Collection Policy/Security	0h 4000 0001h	Page identification Policy access tag
6000 0006h	Collection Error Recovery	0h 1h 3h 4h	Page identification Collection damage summary Last damaged object data time Last damaged attributes time

Table C.1 — Numerical order attributes defined by this standard (part 5 of 6)

Page Number	Page Name	Attribute Number	Attribute
9000 0000h	Root Directory	9000 0000h	"INCITS T10 Root Directory"
		9000 0001h	"INCITS T10 Root Information"
		9000 0002h	"INCITS T10 Root Quotas"
		9000 0003h	"INCITS T10 Root Timestamps"
		9000 0005h	"INCITS T10 Root Policy/Security"
		9000 0006h	"INCITS T10 Root Error Recovery"
9000 0001h	Root Information	0h	Page identification
		3h	OSD System ID
		4h	Vendor identification
		5h	Product identification
		6h	Product model
		7h	Product revision level
		8h	Product serial number
		9h	OSD name
		Ah	Maximum CDB continuation length
		80h	Total capacity
		81h	Used capacity
		83h	Object accessibility
		C0h	Number of partitions
		100h	Clock
		110h	Default isolation method
		111h	Supported isolation methods
		120h	Data atomicity guarantee
		121h	Data atomicity alignment
		122h	Attributes atomicity guarantee
		123h	Data/attributes atomicity multiplier
		1C1h	Maximum snapshots count
		1C2h	Maximum clones count
		1CCh	Maximum branch depth
		200h	Supported object duplication method
		...	" " " "
		2FFh	" " " "
		300h	Supported time of duplication method
		...	" " " "
		30Fh	" " " "
		310h	Support for duplicated object freezing
		0700 0001h	Supported CDB continuation descriptor type
		...	" " " " "
		0700 FFFFh	" " " " "

Table C.1 — Numerical order attributes defined by this standard (part 6 of 6)

Page Number	Page Name	Attribute Number	Attribute
9000 0002h	Root Quotas	0h	Page identification
		1h	Default maximum user object length
		1 0001h	Partition capacity quota
		1 0002h	Partition object count
		1 0081h	Partition collections per user object
		2 0002h	Partition count
9000 0003h	Root Timestamps	0h	Page identification
		2h	Attributes accessed time
		3h	Attributes modified time
		FFFF FFFEh	Timestamp bypass
9000 0005h	Root Policy/Security	0h	Page identification
		1h	Default security method
		2h	Oldest valid nonce limit
		3h	Newest valid nonce limit
		6h	Partition default security method
		7h	Supported security methods
		9h	Adjustable clock
		Ah	Boot epoch
		7FFDh	Master key identifier
		7FFEh	Root key identifier
		8000 0000h	Supported integrity check value algorithm
		...	" " " " "
		8000 000Fh	" " " " "
		8000 0010h	Supported DH group
		...	" " "
		8000 001Fh	" " "
9000 0006h	Root Error Recovery	0h	Page identification
		1h	Root damage summary
		2h	Contained objects damage summary
		3h	Last damaged object data time
		4h	Last damaged object attributes time
		5h	Last damaged contained object time
		6h	Number of damaged partitions
FFFF FFFEh	Current Command	0h	Page identification
		1h	Response integrity check value
		2h	Object Type
		3h	Partition_ID
		4h	Collection_Object_ID or User_Object_ID
		5h	Starting byte address of append

Annex D

(Informative)

Examples of OSD Operation

D.1 Preparing a device for OSD operation

Before an OSD logical unit may accept and process OSD commands, it needs to be initialized as an OSD logical unit. An application client issues the commands in table D.1 to initialize an OSD logical unit.

Table D.1 — OSD initialization sequence

Action	Parameters	Description
SET MASTER KEY	SEED EXCHANGE, DH_Group, DH_Data	Exchange DH seed
SET MASTER KEY	CHANGE MASTER KEY, DH_Data	Initialize master key
SET KEY	Root, Seed	Initialize root key
SET KEY	Partition, Seed	Initialize partition zero key
SET KEY	Working, Key Version, Seed	Initialize partition zero working key
FORMAT OSD	Length (optional)	Construct OSD control structures
CREATE PARTITION	Partition_ID (optional)	Initialize partition in which user objects may be created
SET KEY	Partition Key, Working Keys	Initialize partition keys

Upon completion of these commands the storage device is an OSD logical unit with security established.

D.2 Example of accessing data on an OSD

File system function is beyond the scope of this standard. In this example a simple PC/UNIX-like file system is used. The file system consists of a single file called son in a single subdirectory called father. In PC/UNIX file system notation, the file name would be written as:

/father/son

Table D.2 lists the sequence of OSD commands that may result in the file system being created. It is assumed that the OSD logical unit and the partition are known and that the application client holds a valid capability for each object accessed including an integrity check value.

Table D.2 — OSD command sequence for creating a file

Step	Service Action	Partition_ID	User_Object_ID	Discussion
1	READ	n	fsroot dir	Make sure directory father does not already exist.
2	CREATE	n		Returns User_Object_ID (s), to hold file son.
3	CREATE	n		Returns User_Object_ID (f), to hold directory father.
4	WRITE	n	s	Write contents of file son. ^a
5	WRITE	n	f	Write contents of directory father. ^a
6	WRITE	n	fsroot dir	Root directory revised to contain directory father.
^a More than one WRITE command may be used.				

The CREATE AND WRITE command is capable of transferring data to the newly created object. As is shown in table D.3, separate WRITES are not needed to place data in file son or directory father when this option is used.

Table D.3 — OSD command sequence using CREATE AND WRITE

Step	Service Action	Partition_ID	User_Object_ID	Discussion
1	READ	n	fsroot dir	Make sure directory father does not already exist.
2	CREATE AND WRITE	n		Returns User_Object_ID (s), that holds file son.
3	CREATE AND WRITE	n		Returns User_Object_ID (f), that holds directory father.
4	WRITE	n	fsroot dir	Root directory revised to contain directory father.